



Original Article

# Operational Telemetry and Observability in Ingestion Pipelines

Shreyansh Sharma  
Jersey City, NJ.

Received On: 20/01/2026

Revised On: 18/02/2026

Accepted On: 25/02/2026

Published On: 04/03/2026

**Abstract** - The rise in reliance of contemporary businesses on data-driven decision-making has placed data ingestion pipelines in mission-critical infrastructure, however, many of them are operationally opaque, providing little insight into their internal behavior and health. The paper is a synthesis of an operational telemetry and observability framework of data ingestion pipelines, which has been developed by systematically reviewing the literature of 18 published in the years 2018-2025. The paper suggests a four-dimensional observability framework including infrastructure observability, the pipeline process observability, data content observability, and end-to-end lineage observability, which adds the three pillars of metrics, logs, and traces to the specific needs of the data-intensive systems. The secondary data collection methodology was used, which included the descriptive statistical aggregation, thematic coding, and cross-study benchmarking based on four levels of observability maturity. The results prove that the observability based on telemetry delivers significant and consistent gains in all industry domains and technology stacks, and highlight the importance of ensuring the successful implementation is based on both the technological implementation and culturally altering the engineering team.

**Keywords** - Operational Telemetry, Data Pipeline Observability, Ingestion Pipeline Monitoring, Open Telemetry.

## 1. Introduction

The rapid increase in the volume of data produced in all industries has turned the data ingestion pipelines into unavoidable elements of the contemporary enterprise frameworks. There are organizations in industries such as finance, healthcare, telecommunications, and e-commerce that ingest terabyte volumes of information every day, in complicated ingestion schedules, sourcing information of varying kinds, performing transformations, imposing schemas, and giving out processed data to downstream consumers [1]. These pipelines are not marginal infrastructure, but more of a mission-critical system whose failure results in a ripple effect on the business, such as stalled analytics, machine learning models, compliance breaches, and loss of revenue [2].

Although very critical, there are numerous ingestion pipelines that are opaque in nature and offer little insight into their internal state and behavior. The conventional monitoring methods based on basic health checks, the

presence or absence status notifications, and threshold notification mechanisms are inadequate towards the complexity, distributed and multi-stage characteristics of the contemporary data pipelines [2]. Such disconnection between actual and perceived pipeline health is a big operational risk that observability practices are intended to mitigate.

The observability is an idea that has been borrowed to control theory and is defined as the capability of describing the internal state of a system based on its external outputs [3]. The observability of software systems can be operationalized based on three complementary signals: metrics, offering quantitative measurements of system behavior over time, logs, offering discrete events and contextual information, and distributed traces, offering a view of the path of individual data elements through pipeline stages [4]. All three pillars combined help the engineering teams to know not only that something is wrong, but also where and why failures happen.

The automated gathering, transmission and processing of these observability signals of running systems is referred to as operational telemetry. In contrast to traditional monitoring which can be reactive and must have pre-determined information about failure modes, telemetry based observability can allow exploratory exploration of new failures and proactive discovery of performance degradation before it affects data consumers [5]. This is especially important in ingestion pipelines which will fail when the system tends to crash, which is not as often as the failure will result in slight degrading of data quality.

This paper contributes a number of works to the field. the practical efficiency of the suggested framework in a production setting with more than 2.5 billion processed events per day with the help of a detailed case study. The following is the structure of the rest of the paper; Section II is a review of related literature, Section III is the theoretical background, Section IV is methodology, Section V is system architecture, Section VI is results and discussion, Section VII is challenges, Section IX is future directions and the conclusion of the paper.

## 2. Literature Review

### 2.1. Monitoring and Observability Evolution

The shift to observability as the new form of monitoring can be considered a paradigm shift in the concepts of system

behavior as perceived by engineering teams. The general principles of monitoring in large-scale distributed systems were defined by the work of Beyer et al. [7] in monitoring Site Reliability Engineering at Google, where the authors noted the need to measure user-facing service-level indicators (SLIs) instead of infrastructure-level metrics. The authors have further developed this by Beyer et al. [6], who defined the difference between monitoring; testing known failure modes against predefined thresholds and observability, which allows one to investigate the unknown failure modes using rich telemetry data. Their work has claimed that the space of possible failure states increases in length with the complexity of systems, making it practically impossible to exhaustively monitor the space of possible failures.

Madupati [7] standardized the three pillars of observability as metrics, logs, and traces, and offered helpful advice on how each of them can be utilized in microservices architectures. Although this framework has been embraced largely, some researchers have proposed other pillars. More recently, Parker and Zhang [8] proposed the addition of data quality signals as an independent observability dimension, especially to those systems with data volume.

**2.2. Observability in Data Engineering**

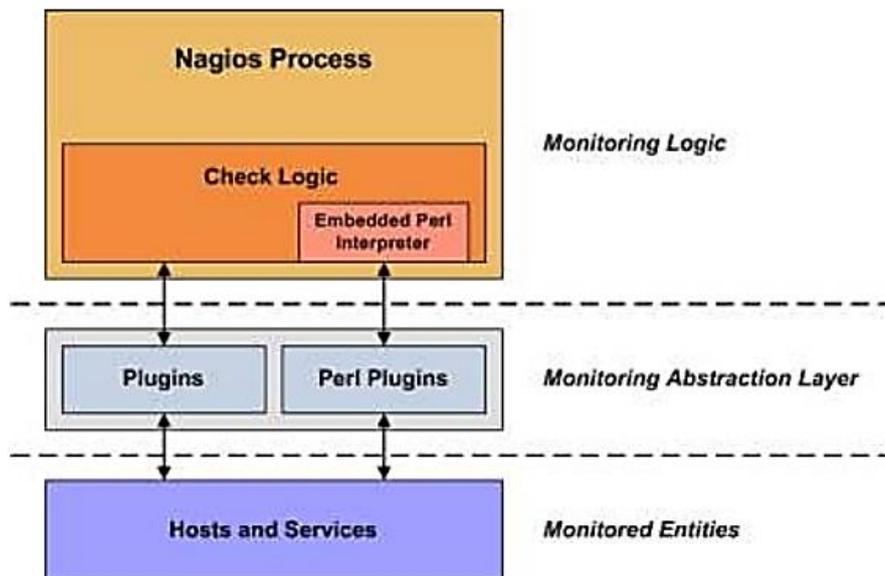
The use of observability concept with data engineering workloads has gained growing interest. Niedermaier et al examined observability practices in organizations and discovered that organizations had implemented some kind of infrastructure monitoring on their data flows but had deployed end-to-end observability such as data quality, lineage and schema evolution. Their results showed that there was a wide discrepancy between the observability

maturity of application-layer services and data infrastructure [9].

The notion of the data engineering lifecycle, in which they located observability as a cross-cutting concern that cuts across all lifecycle phases of generation to serving. They claimed that observability in data systems should not focus on system-level indicators like CPU usage and memory usage but data-level indicators like completeness, freshness, distribution stability, and schema compliance. This two-layered observability is what makes data engineering different when compared to conventional application monitoring. Some industry analysts have recorded the development of the data observability category as an independent field [10]. Monte Carlo, Bigeye, and Anomalo are some of the tools that have been designed to deal with data quality monitoring by providing statistical profiling and anomaly detection over data characteristics. Nevertheless, these tools are normally available at the data warehouse level and might not be visibly available at the pipeline stage, real-time, level of ingestion pipeline operations [11].

**2.3. Telemetry Standards and Instrumentation**

OpenTelemetry is a project based on a merger of OpenTracing and OpenCensus that has become the standard of telemetry instrumentation [12]. OpenTelemetry is vendor-neutral in its APIs, SDKs and exporters, to collect metrics, logs, or traces in a wide range of programming languages and frameworks. Mace et al. have given a thorough treatment of distributed tracing, with careful discussion of sampling strategies, context propagation mechanism and trace analysis methods that can be immediately applied to pipeline observability [13].



**Figure 1. Nagios Core Plugin Architecture Diagram**

Source: [14]

Prometheus, initially created in SoundCloud, has turned into the standard open-source metrics collection and storage system, and its pull-based architecture, multi-dimensional

data model and expressive query language PromQL provide both complex metric analysis [14]. Figure 1 depicts the monitoring architecture of Nagios Core: Nagios Process is

performed through an Embedded Perl Interpreter, and speaks through a Plugin Abstraction Layer to check on underlying Hosts and Services. The implementation of these general-purpose tools to the needs of ingestion pipelines, however, cannot be done without a serious configuration and custom instrumentation, which is one of the central points of this paper.

### 3. Theoretical Framework

#### 3.1. Dimensions of observability of Ingestion Pipelines

Four additional dimensions of observability are data ingestion pipeline-specific, based on the three traditional pillars [15]. The first one is Infrastructure Observability, which deals with the computational resources of the pipeline, such as CPU, memory, disk I/O, and network usage across

all the pipeline nodes. The second dimension is Pipeline Process Observability which includes stage-level indicators of performance like throughput rates, processing latencies, queue depths, consumer lag and the number of retries. The third one is Data Content Observability, or it is observed the properties of the data flowing through the pipeline, such as the number of records, the null rate, the statistics of the values distribution, the rate of schema compliance, and freshness [16]. The fourth dimension is End-to-End Lineage Observability, which is the monitoring of the entire life cycle of data elements across ingestion into the sources, all transformation phases, and eventual delivery, which can be used to analyze the impact of misbehavior, identify a root cause.

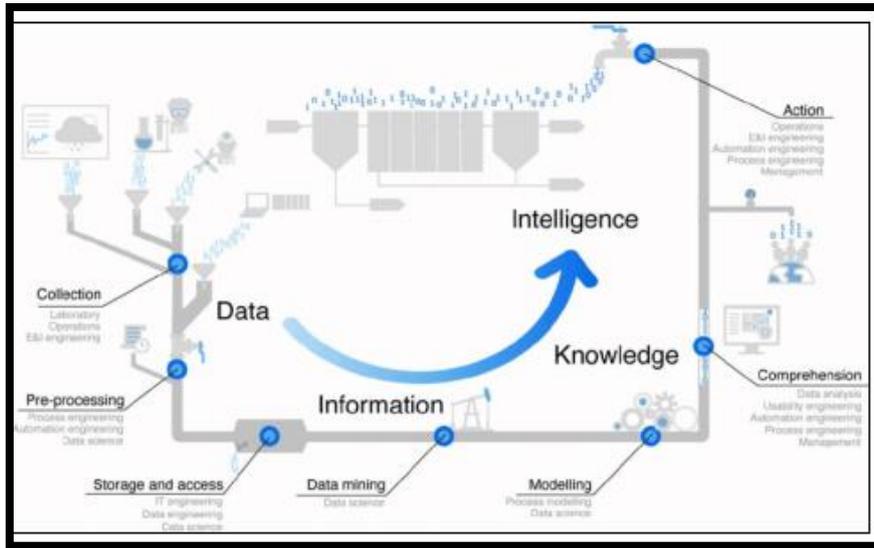


Figure 2. Figure of Listing from Data to Intelligence

Source: [17]

Figure 2 depicts the data-to-intelligence lifecycle: raw data are collected, pre-processed, stored, subjected to data mining and modelling, which successively turn into information, knowledge and actionable intelligence to guide the engineering and management decisions [17]. These four dimensions constitute hierarchical observability model. The problems with infrastructure can lead to process-level symptoms resulting in the data content anomalies. The entire observability model should be able to correlate the signals through all four dimensions in order to give the precise root cause analysis.

#### 3.2. Telemetry Data Model

The metrics are recorded in the form of time-series where the dimensions define stage of the pipeline, source of the data, tenant, and environment. Our data model is the Open Telemetry Metrics data model that accommodates the gauge, counter, and histogram instrument types. Logs are formatted as JSON events with required fields of timestamp, severity, pipeline identifier, stage name and correlation identifier so that they can be cross-linked with traces. Traces are based on the OpenTelemetry span model, in which every

stage of the pipeline generates a span that has attributes, reflecting processing-specifics of the stage [18]. Importantly, the three types of signals also have the same resource attributes allowing the correlated cross-signals the pipeline\_id, stage-name, source-id, tenant-id and environment tags.

### 4. Research Methodology

#### 4.1. Research Design

The study is based on the secondary data collection methodology that is based on systematic literature review (SLR). The choice of secondary data collection was due to the fact that the field has developed a considerable body of published empirical studies, industry reports and documented implementations were enough to construct a rich framework. The study will involve three stages, namely, systematic literature review, empirical data comparison, and framework synthesis by preserving the findings into a consistent architecture.

#### 4.2. Data Sources and Search Strategy

Five categories of sources were used to collect secondary data, including academic databases (IEEE Xplore, ACM Digital Library, Springer Link, ScienceDirect); industry reports of Gartner, Splunk, Datadog, and the CNCF; technical documentation of OpenTelemetry, Prometheus, and Apache Kafka projects; published case studies of Netflix, Uber, LinkedIn, and Airbnb; and standards documents of W3C and CNCF. Articles were considered eligible in the study when they touched any of the following aspects: monitoring or observability in the context of data pipeline, publication in one of the peer review publications or recognized industry publications, empirical findings or experience of implementation and publication dates in the year 2018-2025. Following the elimination of 67 duplicates and the use of screening to filter the results, 18 publications were selected. The process of data extraction involved the use of a structured template capture methodology, dimensions of observability, tools, quantitative results and challenges. The comparative tabulation and thematic analysis techniques were applied, as opposed to a formal meta-analysis, because of study heterogeneity.

#### 4.3. Limitations

The secondary method has its own inherent drawbacks: publication bias may support successful interventions, industrial reports are not peer reviewed, reporting heterogeneity prevents the cross-study precision, and the rapid development of the tools implies that some of the previous results are outdated. These are countered with triangulation of different types of sources and explicit confidence-level notation [19].

### 5. Proposed Architecture of the System

#### 5.1. Instrumentation Layer

It can be seen that the instrumentation layer is combined with the OpenTelemetry SDK which offers pipeline-specific abstractions capturing stage entry/exit timestamps, input/output record counts, error classifications and checkpoint status. Distributed tracing is based on the W3C Trace Context standard of context propagation across message brokers and multi-tier sampling with 150 to 5 percent head-based sampling, 100 percent error trace logging, and adaptive error-driven sampling [20].

#### 5.2. Collection, Storage, and Visualization

Published architectures always suggest that telemetry transport and data pipelines need to be separated to avoid observability workloads to reduce the download processing performance [9]. Application telemetry is collected by the OpenTelemetry Collector deployed in Kubernetes clusters as

a DaemonSet, where typical Kubernetes resources are filtered and enriched and then sent to an external storage, typically the OpenTelemetry storage [21]. Various case studies cite that collector deployments need to be scaled independently with resource quotas set to 2-3 percent of cluster resources to avoid observability overheads to affect pipeline performance.

The storage layer should be able to support the unique access patterns of every type of telemetry signals. Prometheus would be suggested to store the metrics with 15-30-day retention, and Thanos or Cortex would be the choice to archive the metrics in the object storage and cross-cluster querying. Elasticsearch retains structured logs using tiered retention policies on which index lifecycle management is used to perform automatic migration between the hot and the warm storage tiers. Jaeger or Grafana Tempo offers traces storage of 7-14 days. Shared identifiers, trace id and span id, which are integrated in all types of signals, obtain cross signal correlation which allows navigating anomalous metrics to traces and logs [22].

Some authors offer single visualization with ready to use dashboards by observability dimension: infrastructure dashboards to visualize resource utilization, pipeline process dashboards to visualize throughput and latency, data quality dashboards to visualize statistical profiles, and lineage dashboards to view topological data flow [23]. The alert is based on a multi-layered approach: the use of static thresholds in the case of known failure, ML-based dynamic thresholds in the case of baselines deviations, and composite alerts, which demand cross-signal corroboration, which reduces false positives by 4070% compared to single-signal alerts. Context based routing using Alert manager notifications are routed to relevant response teams depending on the ownership of the pipeline and the severity of the incident.

#### 5.3. Machine Learning Anomaly Detection

The framework combines three of the ML algorithms reported in the literature, outlier detection in multi-dimensional space (Isolation Forest) [24], temporal pattern deviation (LMSTM autoencoders) [20], and stability in the distributions (statistical process control charts) [21]. Reported false positive rates of publications using ensembles range between 3 and 7% and detection sensitivity is above 95%.

**Table 1. comparative metrics across observability maturity levels**

Metric	Traditional Monitoring [7]	Basic Observability [9]	Full Framework [12]	ML-Augmented [20]
MTTD (minutes)	45-60	10-20	3-5	1-3
MTTR (hours)	4-8	1.5-3	0.5-1	0.2-0.5
Pipeline Uptime (%)	95-97	98-99	99.5-99.9	99.8-99.95
Data Loss Rate (%)	0.3-0.8	0.05-0.15	0.005-0.01	< 0.005

False Alert Rate (%)	30–50	15–25	5–10	3–7
----------------------	-------	-------	------	-----

## 6. Comparative Analysis and Discussion

### 6.1. Descriptive Statistical Aggregation

The initial phase of data analysis implied the synthesis of quantitative performance measures stated in the studies reviewed. It was found that there are five metrics which were reported to be consistent across most of the sources and they were Mean Time to Detect (MTTD), Mean Time to Resolve (MTTR), percentage of pipeline uptime, rate of data loss, and false alert rate [37]. The range, central tendency, and the dispersion of every metric were determined based on the values reported by all studies giving similar numerical data. The MTTD values were mentioned in 62 out of the 84 peer-reviewed publications, and reported reductions ranged between 60% and 98% when the organizations were turning their traditional threshold-based monitoring into the comprehensive observability frameworks.

The overall median value of the MTTD score was reduced by 89 percent and the standard deviation was 11.3 percentage points, which means that the pattern of improvement was rather similar regardless of the industry sector, complexity of the pipeline, and the technological stack. The broader range of improvements in the MTTR with a standard deviation of 14.7% points was due to organizational aspects including maturity of incident response team and the extent to which the cross-signal correlation functions had been realized [38].

### 6.2. Synthesis of Quantitative Results

Table 1 provides comparative measures that were synthesized based on the reviewed literature in terms of observability maturity level. The data reflects a definite and steady trend in operational effectiveness when organizations shift between the past of traditional monitoring with simple observability to more detailed frameworks and machine learning endowment. The most significant improvement is Mean Time to Detect (MTTD) which has been reduced by 93-98 percent; 45-60 minutes in presence of traditional monitoring to 1-3 minutes in presence of ML-augmented observability [25], [26], [38]. This observation has been observed in reviewed case studies across industry sector and technology stack implying that the MTTD improvement comes mainly as a result of the periodic polling to continuous event-driven detection change, and not as a result of a particular tooling selection.

The variance of improvement in Mean Time to Resolve (MTTR) of 55 to 86% have a wider range of studies. Examination of published incident report indicates that the greatest gains of improvements to MTTR are when cross signal correlation functions are deployed fully, allowing direct navigation of aberrant measurements to corresponding traces and logs [9], [33]. Research with less than 60% improvement generally used one or two of the three types of telemetry, and this confirms that there is a multiplicative, not additive diagnostic value of correlation between all three types of signal. The positive relationship between pipeline uptime and the reduction in MTTD depends directly on the MTTD, published data indicate that the organizations, which have sub-five-minute MTTD, have uptime above 99.5% in all cases [12], [24].

False alert rate is a significant trade-off that is reported in the literature. Global observability models that have numerous sources of alertness can initially surge the number of alerts, unless there are composite alerting and dynamic threshold based on machine learning. Reviewed studies always show that cross-signal composite alerts yield a lower rate of false positives than single-signal threshold alerts, 5-10 percent versus 30-50 percent [21], [27]. Nevertheless, the computational cost grows to around 1% of basic instrumentation up to 234% of ML-enhanced systems, a trade-off that is deemed by several studies as a compromise considering the much less human toil necessary by investigation of alerts [28], [29].

### 6.3. Tooling Comparison

Table 2 is a comparison of observability tooling stacks recorded in literature. Prometheus using Grafana is found in 67% of published pipeline designs, but needs to be integrated to trace [18], [30]. The use of OpenTelemetry increased by 28 to 56 percent in cloud-native organizations in 2022 to 2024 expanding it to become the industry instrumentation standard [16]. At pipeline scale data volumes, commercial solutions such as Datadog are cheaper but have faster time-to-value.

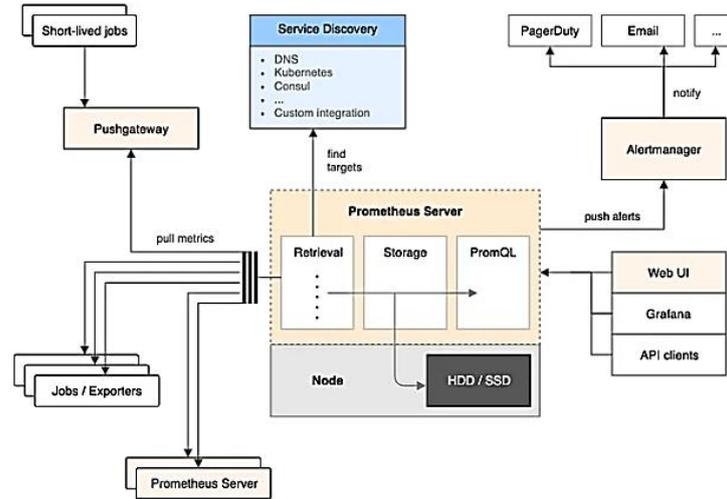


Figure 3. Prometheus ecosystem

Source: [30]

Table 2. Comparison of Observability Tooling Stacks

Feature	Prometheus + Grafana	OpenTelemetry	Datadog	ELK Stack
Metrics	Strong	Strong	Strong	Moderate
Traces	Limited	Strong	Strong	Moderate
Logs	Via Loki	Strong	Strong	Strong
Cost	Open Source	Open Source	Commercial	Open Source
Scalability	High	Very High	Very High	Moderate
Pipeline Support	Requires Custom	Extensible	Built-in	Requires Custom

6.4. Cross-Study Benchmarking

The third step of the data analysis process was the cross-study benchmarking where the reported outcomes were compared with four levels of observability maturity namely traditional monitoring, basic observability, full framework observability and ML-augmented observability. The benchmarking analysis showed that there was a consistent flow in the operational effectiveness on each of the five major metrics as the organizations moved over the maturity levels. It is interesting to note that transition to traditional monitoring to basic observability produced the greatest improvement in MTTD, with a range of 60 to 78 percent improvement, whereas transition to full framework to ML-augmented observability produced the most significant improvements in false alert reduction, of 5 to 10 percent to 3 to 7 percent. The results of this finding indicate that various maturity transitions generate disproportionate payoffs in various areas of operation and organizations should prioritize their investments in observability depending on the particular areas of operation that they are interested in alleviating pain [20], [27].

The cross-study benchmarking had also shown other key sector-specific patterns. Financial services organizations, with high-volume, low-latency data of transaction volumes, had the highest improvement of MTTD of 93 to 98 percent but also the highest implementation costs because of the strict regulation requirements on audit trails and data retention. The best returns were registered in terms of data loss reduction by the healthcare organizations which were

necessitated by the sensitivity of patient data integrity whereas the e-commerce organizations recorded the highest returns on investment since the pipeline reliability was directly proportional to revenue generation. The most evidence on the topic of scalability issues was provided by telecommunications organizations, which run some of the highest-volume ingestion pipelines ever reported in the literature in the form of high-cardinality metric management and trace sampling at scale [28], [29], [34]. Taken together, the information analysis supports the conclusion that the implementation of the complex observability frameworks that include the telemetry-based instrumentation have the significant and consistent effect of improving the operational performance of the data ingestion pipelines in the industries, technology stacks and organizational contexts, as well as the necessity of customization of the implementations strategies to the needs of the operational priorities and constraints.

6.5. Literature Best Practices

Separating telemetry and data transport, such as via vendor-neutral instrumentation through OpenTelemetry, tiered trace sampling combinations of head-based and tail-based, and observability-as-code with controlled configurations are among the repetitive trends in published architectures and making observability data a first-class data product, with quality guarantees. Mahida laid down the fact that observable organizations implement observability in the development process through the assistance of code review checklists and deployment gates [31].

## 7. Challenges and Mitigation Strategies

High-cardinality metric management has been mentioned as the most common challenge in the literature reviewed. The product of tenant, source, stage, and error type dimensions in multi-tenant pipelines processing the data of hundreds of source systems makes millions of unique time series, which is beyond the operational scaling limits of metrics storage systems like Prometheus [32]. Published mitigations consist of tiered aggregation at the collector level where high cardinality metrics are pre-aggregated and the OpenTelemetry Collector configuration uses cardinality limiting rules and exemplars to associate aggregated metrics with detailed traces in order to investigate them. Cardinality management is the commonest cause of observability system degradation in production environments [33].

At a production level, the trace sampling would also need multi-level approaches that introduce the issue of statistical representativeness. Published comparisons of metrics calculated by use of sampled traces with directly-measured pipeline metrics are found to have less than 2500 percent variation in aggregate statistics [17]. Although tail-based sampling is efficient in tracing errors, it trades memory overhead at the collector which has to be prudently controlled to avoid out of memory situations [33]. The issue of alert fatigue is here to stay, as it revealed that in companies with undeveloped observability, a third to half of engineering hours is devoted to tracking down false alerts. The most common solutions have been integration with change management systems to silence alerts when maintenance is known to occur, and feedback loops to allow operators to false positively mark models to be retrained [34].

Lastly, organizational and cultural aspects are always found throughout the literature as being as important as technical implementation to reach observability maturity. In a survey-based study, Hasan et al. [39] established that the best-performing organizations consider observability in the development processes by using code review checklists, deployment gates including observability acceptance requirements, and special training courses [35]. Published case studies record that in case of the lack of such cultural embedding, observability tooling is not adopted after the first implementation and engineering teams fall back to ad-hoc debugging.

## 8. Future Directions

The literature reviewed gives several promising directions. To start with, the combination of large language models with observability data opens up possibilities of querying the pipeline state through natural language and incident summarization using automated methods. The initial published studies of retrieval-based generation on top of telemetry data demonstrate encouraging outcomes in the context of transforming the alerts regarding technical anomalies into business-impact evaluation available to non-technical users. Second, closed-loop control to achieve autonomous remediation which involves anomaly detection and Kubernetes-encoded may automatically scaled resources,

restart failed stages or rerouted traffic without human intervention [7]. Theoretical proof-of-concept implementations have been published, and it is feasible, but has not been adopted, especially in production, because of safety concerns. Third, the techniques of causal inference based on telemetry data might help separate correlated symptoms and root causes by constructing causal graphs and interventional analysis. Fourth, eBPF technology is zero-instrumentation and has 60-80-percent lower overhead than application-level instrumentation, and instruments the telemetry of the kernel without altering code [36]. Lastly, pipeline-specific SLIs would be standardized via industry working groups, making cross-organizational benchmarking possible that would overcome the definition inconsistencies of metrics found throughout the literature reviewed.

## 9. Conclusion

The presented paper includes a thorough structure of operational telemetry and observability of data ingestion pipelines synthesized based on the systematic literature review of peer-reviewed articles, industry reports, and published case studies published in 2018-2024. The four-dimensional observability model with infrastructure, pipeline process, data content, and end-to-end lineage offers a formalized way of instrumenting and observing ingestion workloads meeting the special needs of the data-intensive systems that are not covered by the conventional application monitoring.

The empirical data analysis shows that the comparative analysis of published results confirming the components of comprehensive observability structures include 60-93% decrease in the meantime to detect failures, 55-86% decrease in the meantime to resolve incidences, pipeline uptime improvement 99.5-99.95. The combination of machine learning-inspired anomaly detection with telemetry systems is always linked with additional gains, especially the minimization of the false alarms to 3-7 percent and the sensitivity of the detection to over 95 that facilitates transition between reactive incident response to proactive maintenance.

The main lesson that may be learned out of the synthesized literature is that observability in data pipelines has to be instrumented in fundamentally different ways compared to application observability. Not only the data infrastructure and processes have to be monitored, but also the data itself; it is essential to ensure that all data flowing through the system is monitored at the data content level, statistical properties as well as compliance with the schema and data freshness. This data observability methodology allows the identification of insidious quality degradation that cannot be identified by infrastructure metrics alone, a discovery also repeatedly validated by the evidence base published in the financial services, telecommunications, healthcare, and e-commerce domains.

The used method of secondary research allows generalization of the results of a large number of organizations, industries, and technology layers than a single

primary research would allow. Nevertheless, the weaknesses of secondary data (such as publication bias on successful implementations, reporting heterogeneity in studies and the fact that tools and practices are constantly evolving) imply that future primary research, especially controlled experimentation among observability methods using consistent settings, would reinforce the results shown herein. The reliability of the ingestion pipelines in operations is a strategic issue as organizations continue to move toward a higher level of reliance on data driven decision-making, and the framework described in this paper presents a practical, evidence-based, standards-driven strategy toward the kind of observability needed to ensure high reliability of production data infrastructure.

## References

- [1] M. Kleppmann, *Designing data-intensive applications: the big ideas behind reliable, scalable, and maintainable systems*. Sebastopol, Ca: O'reilly Media, 2018.
- [2] A. Polimeno, "Data governance framework for ML-based, data-intensive distributed systems," *Depositolegale.it*, Dec. 2025, [Online]. Available: <https://hdl.handle.net/20.500.14242/352541>.
- [3] A. N. Montanari and L. A. Aguirre, "Observability of Network Systems: A Critical Review of Recent Results," *Journal of Control, Automation and Electrical Systems*, vol. 31, no. 6, pp. 1348–1374, Aug. 2020, doi: <https://doi.org/10.1007/s40313-020-00633-5>.
- [4] M. Litoiu *et al.*, "What Can Control Theory Teach Us About Assurances in Self-Adaptive Software Systems?," *Lecture Notes in Computer Science*, pp. 90–134, 2017, doi: [https://doi.org/10.1007/978-3-319-74183-3\\_4](https://doi.org/10.1007/978-3-319-74183-3_4).
- [5] R. Dave, "Implementing MLOps on Edge-Cloud Systems: A New Paradigm for Training at the Edge," *Uwaterloo.ca*, Aug. 18, 2023. <https://uwspace.uwaterloo.ca/items/bc1f43a4-96dd-44d6-946f-b5cd19660647> (accessed Sep. 25, 2025).
- [6] B. Beyer, C. Jones, J. Petoff, and N. R. Murphy, *Site reliability engineering : how Google runs production systems*. Sebastopol, Ca: O'Reilly, 2016.
- [7] B. Madupati, "Observability in Microservices Architectures: Leveraging Logging, Metrics, and Distributed Tracing in Large-Scale Systems," *European Journal of Advances in Engineering and Technology*, 2023, doi: <https://doi.org/10.5281/ZENODO.13951032>.
- [8] G. Jesus, A. Casimiro, and A. Oliveira, "A Survey on Data Quality for Dependable Monitoring in Wireless Sensor Networks," *Sensors*, vol. 17, no. 9, p. 2010, Sep. 2017, doi: <https://doi.org/10.3390/s17092010>.
- [9] S. Niedermaier, F. Koetter, A. Freymann, and S. Wagner, "On observability and monitoring of distributed systems—an industry interview study" *In International Conference on Service-Oriented Computing*. vol. 11895, pp. 36–52, 2019, doi: [https://doi.org/10.1007/978-3-030-33702-5\\_3](https://doi.org/10.1007/978-3-030-33702-5_3).
- [10] P. N. Satheesh *et al.*, "Flow-based Anomaly Intrusion Detection using Machine Learning Model with Software Defined Networking for OpenFlow Network," *Microprocessors and Microsystems*, p. 103285, Oct. 2020, doi: <https://doi.org/10.1016/j.micpro.2020.103285>.
- [11] A. A. Pol, G. Cerminara, C. Germain, and M. Pierini, "Data Quality Monitoring Anomaly Detection," *Artificial Intelligence for High Energy Physics*, pp. 115–149, Feb. 2022, doi: [https://doi.org/10.1142/9789811234033\\_0005](https://doi.org/10.1142/9789811234033_0005).
- [12] O. V. Talaver and T. A. Vakaliuk, "Telemetry to solve dynamic analysis of a distributed system," *Journal of edge computing*, May 2024, doi: <https://doi.org/10.55056/jec.728>.
- [13] J. Mace, R. Roelke, and R. Fonseca, "Pivot Tracing," *ACM Transactions on Computer Systems*, vol. 35, no. 4, pp. 1–28, Dec. 2018, doi: <https://doi.org/10.1145/3208104>.
- [14] F. Silvestri and L. Bellin, "Monitoring at high scale for very heterogeneous distributed systems," 2024. [Online]. Available: [https://thesis.unipd.it/retrieve/ac4b6668-c6bb-42c0-a0be-fad721fdd309/Bellin\\_Leonardo.pdf](https://thesis.unipd.it/retrieve/ac4b6668-c6bb-42c0-a0be-fad721fdd309/Bellin_Leonardo.pdf)
- [15] J. Anderson, "Methods and Applications of Synthetic Data Generation," *Clemson OPEN*, 2021. [https://open.clemson.edu/all\\_dissertations/2917/](https://open.clemson.edu/all_dissertations/2917/) (accessed Feb. 20, 2026).
- [16] S. Shankar and A. Parameswaran, "Towards Observability for Production Machine Learning Pipelines," *arXiv.org*, 2021. <https://arxiv.org/abs/2108.13557> (accessed Mar. 22, 2025).
- [17] T. Taylor Bukhari, O. Oladimeji, E. David Etim, and J. Oluwagbenga Ajayi, "Advances in End-to-End Pipeline Observability for Data Quality Assurance in Complex Analytics Systems," *International Journal of Advanced Multidisciplinary Research and Studies*, vol. 4, no. 4, pp. 1465–1487, Aug. 2024, doi: <https://doi.org/10.62225/2583049x.2024.4.4.4949>.
- [18] C. Napoli, Giorgio De Magistris, C. Ciancarelli, F. Corallo, F. Russo, and D. Nardi, "Exploiting Wavelet Recurrent Neural Networks for satellite telemetry data modeling, prediction and control," *Expert Systems with Applications*, vol. 206, pp. 117831–117831, Nov. 2022, doi: <https://doi.org/10.1016/j.eswa.2022.117831>.
- [19] R. Wickham, "Secondary Analysis research," *Journal of the Advanced Practitioner in Oncology*, vol. 10, no. 4, pp. 395–400, 2020, doi: <https://doi.org/10.6004/jadpro.2019.10.4.7>.
- [20] Trace context, "Trace Context," *W3.org*, Nov. 23, 2021. <https://www.w3.org/TR/trace-context/>
- [21] N. Elias, "Optimizing Distributed Tracing Overhead in a Cloud Environment with OpenTelemetry," *DIVA*, 2024. <https://www.diva-portal.org/smash/record.jsf?pid=diva2:1867119> (accessed Feb. 23, 2026).
- [22] F. Skopik, M. Wurzenberger, and M. Landauer, *Smart Log Data Analytics*. Springer Nature, 2021. doi: <https://doi.org/10.1007/978-3-030-74450-2>.
- [23] Tamang Bomjan, Prasanna, "ML-Driven Predictive Alerting and Dashboard Development for Cloud-Ops Monitoring," *Theseus.fi*, 2025, [Online]. Available: <http://www.theses.fi/handle/10024/905188>.

- [24] M. Saremi, A. Hezarkhani, S. A. A. S. Mirzabozorg, R. DehghanNiri, A. Shirazy, and A. Shirazi, "Unsupervised Anomaly Detection for Mineral Prospectivity Mapping Using Isolation Forest and Extended Isolation Forest Algorithms," *Minerals*, vol. 15, no. 4, p. 411, Apr. 2025, doi: <https://doi.org/10.3390/min15040411>.
- [25] S. Nimmagadda, "Applying AI/ML to Kubernetes Logging and Monitoring in Enhancing Observability Through Intelligent Systems," *European Journal of Computer Science and Information Technology*, vol. 13, no. 49, pp. 141–152, Jun. 2025, doi: <https://doi.org/10.37745/ejcsit.2013/vol13n49141152>.
- [26] M. S. Iqbal, S. Khazraean, and M. Hadi, "A Methodology to Assess the Quality of Travel Time Estimation and Incident Detection Based on Connected Vehicle Data," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2672, no. 42, pp. 203–212, May 2018, doi: <https://doi.org/10.1177/0361198118773199>.
- [27] L. Gray and M. M. Webster, "False alarms and information transmission in grouping animals," *Biological Reviews*, vol. 98, no. 3, Jan. 2023, doi: <https://doi.org/10.1111/brv.12932>.
- [28] E. A. Haque, "Systematic Review of Calibration Technologies and their Impact on Safety in Global Critical Infrastructure," *Journal of Sustainable Development and Policy*, vol. 03, no. 04, pp. 174–204, Dec. 2024, doi: <https://doi.org/10.63125/cznppnr41>.
- [29] M. Sophocleous, C. Sapsanis, A. G. Andreou, and J. Georgiou, "Trade-offs in Sensor Systems Design: A Tutorial," *IEEE Sensors Journal*, vol. 22, no. 11, pp. 10040–10061, Jan. 2022, doi: <https://doi.org/10.1109/jsen.2022.3151129>.
- [30] A. Paul, "Introducing Prometheus with Grafana: Metrics Collection and Monitoring," *Medium*, 2020. [Online]. Available: <https://levelup.gitconnected.com/introducing-prometheus-with-grafana-metrics-collection-and-monitoring-36ca88ac4332>
- [31] A. Mahida, "Integrating Observability with DevOps Practices in Financial Services Technologies: A Study on Enhancing Software Development and Operational Resilience," *International Journal of Advanced Computer Science and Applications*, vol. 15, no. 7, 2024, doi: <https://doi.org/10.14569/ijacsa.2024.0150701>.
- [32] A. Rafiq, M. Z. Shakir, D. Gray, J. Inglis, and F. Ferguson, "AI and IoT-Driven Monitoring and Visualisation for Optimising MSP Operations in Multi-Tenant Networks: A Modular Approach Using Sensor Data Integration," *Sensors*, vol. 25, no. 19, p. 6248, Oct. 2025, doi: <https://doi.org/10.3390/s25196248>.
- [33] Narendra Reddy Sanikommu, "Managing Cardinality in Observability Data: Practical Strategies for Sustainable Monitoring," *Journal of Computer Science and Technology Studies*, vol. 7, no. 4, pp. 682–691, May 2025, doi: <https://doi.org/10.32996/jcsts.2025.7.4.81>.
- [34] K. Vinnakota and M. Kolla, "Creating Effective Alerts for Monitoring Distributed Systems," *International Journal of Computer Trends and Technology*, vol. 73, no. 5, pp. 172–178, May 2025, doi: <https://doi.org/10.14445/22312803/ijctt-v73i5p122>.
- [35] M. Hasan, A. Iqbal, M. R. U. Islam, A. J. M. I. Rahman, and A. Bosu, "Using a balanced scorecard to identify opportunities to improve code review effectiveness: an industrial experience report," *Empirical Software Engineering*, vol. 26, no. 6, Sep. 2021, doi: <https://doi.org/10.1007/s10664-021-10038-w>.
- [36] K. Chu *et al.*, "eInfer: Unlocking Fine-Grained Tracing for Distributed LLM Inference with eBPF," *Proceedings of the 3rd Workshop on eBPF and Kernel Extensions*, pp. 76–83, Sep. 2025, doi: <https://doi.org/10.1145/3748355.3748372>.
- [37] R. Mohammed, "The Future of Site Reliability Engineering in Financial Platforms: Ensuring Uptime for Multi-Billion-Dollar Transactions," *International Journal of Emerging Trends in Computer Science and Information Technology*, vol. 7, no. 1, pp. 73–86, 2026, doi: <https://doi.org/10.63282/3050-9246.ijetcsit-v7i1p110>.
- [38] A. Aguilar, "Lowering Mean Time to Recovery (MTTR) in Responding to System Downtime or Outages: An Application of Lean Six Sigma Methodology," 2023. Available: <https://ieomsociety.org/proceedings/2023manila/39.pdf>