

Observability in Stateful Workloads: Strategies for Monitoring Persistent Services in Dynamic Cloud Environments

Sunil Agarwal

Submitted: 01/08/2025

Revised: 20/08/2025

Accepted: 02/09/2025


Abstract: Stateful workloads are central elements of any contemporary cloud-native design, but their permanence brings them special difficulties to observation. It introduces an elaborate system of tracking such workloads in terms of metric matching, failure propagation modeling, and cross-layer tracing continuity, outlined in this paper. In order to evaluate how useful the telemetry overhead, the anomaly correlation, and the dashboard-based diagnostics are, we conducted experiments spread over PostgreSQL, Kafka and Redis deployments. The conventional observability methods have proved within our means, to explain how the state changes and long-term associations. To minimize the time of diagnosis as well as enhance its reliability, we suggest optimized approaches to the reduction of metrics and unified dashboards. These plans enable DevOps teams to have scalable resilient operation tools.

Keywords: Cloud, Observability, Workloads, Dynamic

I. Introduction

With the evolution of distributed systems to the microservice and the hybrid cloud architecture, observability has developed into one of the fundamentals

of system reliability. Whereas stateless workloads have built-in monitoring streams, stateful workloads (e.g., databases, brokers, and session stores) are more complicated because they behave in a persistent manner and contain complex internal state.



Observability capabilities to track their real-time performance, failure modes and interactions across services requires a more sophisticated set of observability strategies than classic triads of logs, metrics, and traces. The piece explores the telemetry needs of stateful workloads and observes architectural areas of

concentrations in matters of observability and provides workable steps to building resilient, platform-agnostic monitoring solutions in highly dynamic cloud-native software systems.

II. Related Works

Distributed Systems


The current DevOps- and microservices-driven modern cloud-native architectures have also created tremendous challenges in the observability of the systems specifically around the stateful workloads. These workloads (databases, caches and persistent services) are sensitive to tracking context and tend to be very reliable, so it is necessary to examine them carefully, and thus they may be beyond the scope of traditional stateless monitoring.

Advanced observability is required because cloud environments are getting more dynamic and complicated. As a qualitative study of 28 software professionals showed, although DevOps practices enable such desirable qualities as agility and scalability, they also generate interdependencies, which are difficult to follow in a scenario involving distributed systems [1].

Observability can then not only be a technical requirement but also as a strategic force in resilience and system fitness. Failure to organize around practices of observability is common. The research highlighted the lack of consistency in understanding between the management and the developers and the importance of the strategic, cross-watch observability products entrenching the roles and responsibility into the software delivery lifecycle [1].

Observability of a stateful service must be end-to-end: instrumentation, pipeline architecture and actionability. The urgency of the systemic design plainly resonates via the initiative such as the Sieve platform [2], where the attention is focused on the minimization of metric dimensionality to be able to estimate causality. The latter goal of efficiency and insight is tantamount in the context of workload that is persistent in nature whose anomalous signals may be overridden by telemetry noise during the initialization of severe degradations.

Trace Span Gaps by Component



Integrated Approaches

Conventionally, observability has been dealt with as a three-pronged scenery of logs, metrics, and traces. There is however a problem that these stovepipes work in isolation and as such are not adequate to analyse complex stateful interactions. These difficulties are compounded in systems whose context continuity, e.g. session integrity or transaction state, is central to the diagnosis of system failure.

This is one of the positive directions that are the standardization of observability sources to unify them with structured logging and uniform schema layers. In [3], a unification of logging libraries was presented, which has

large capacities to process thousands of events per minute but maintains the format of structured data and thus does not cause developers any cognitive burden.

The testing assured that unification of events and semantic labeling may help in improved diagnostics cross-layer. Similarly, every microservice tracing and analysis is developing. The survey on multiple companies in the industry has shown that most organizations showed the presence of tracing pipelines, but due to insufficient robust analytics levels, such pipelines are not notably useful, such as root cause localization or anomaly correlation [4].

Visualization is still the primary technique, and it is sometimes inadequate at interpreting stateful transitions, or in showing persistent bottlenecks. One more study hypothesized that log aggregation tools (e.g. Fluentd, Elasticsearch) and distributed tracing (e.g. OpenTelemetry, Jaeger) are the building block of next generation observability stack [8]. Through this integration the end-to-end monitoring is possible including user requests to database write latency which is essential in debugging delays in state retention or replication delays.

Event Modeling

Stateful workloads need observability plans that are able to record longitudinal patterns i.e. patterns that stretch across time and between replications. The anomaly detection needs to graduate to continuous and contextual modeling compared to snap shot based detection.

The recent advancements involve the state machine-based model to detect the repugnant anomalies at run-time of the multi-service clouds [5]. In contrast with deep neural networks, state machines are interpretable, which is a very valuable property in persistent systems where the malclassification of anomaly alerts might corrupt live databases with false positive records. The proposed model registers 99.2 percent of balanced accuracy and an F1 score of 0.982 indicating that it serves well in identifying subtle threats on the run like delay propagation and unauthorized access attempts [5].

Eadro is an integrated framework to anomaly detection and root cause localization, which integrates multi source data (KPIs, logs, traces) into context-specific alerts [10]. It fills the gap between two long-discriminated stages; detection and localization. The design of Eadro suits stateful workload requirements, failure tends to be non-linear, diagnosis involves correlating resource contention, pattern of locking and state corruption across nodes.

Another problem is a silent model failure in the machine learning pipelines, delivered as a part of the production. These failures might be hidden to usual monitoring as there are no real-time feedback loopings. Towards the reduction of this impact, [6] suggests a bolt-on observability architecture that assists in the detection, diagnosis, and reaction to ML-driven systems. These systems often touch stateful stores or have registries, and once again the necessity of resilient and cross layer observability should be confirmed.

Platform Optimization

Stateful observability also overlaps with business process monitoring, and cross-platform performance engineering. In contemporary hybrid and multi-knockdown architectures, the presence of enduring services that span

among the entities and outsider platforms is typical, and end-to-end troubleshooting may not be easy to perform.

Business process observability offers a top-level abstraction that can be used in tracking the well-being of workflows, particularly, where cloud-native systems communicate with legacy systems [7]. This model constructs a topology graph which connects Biz KPIs with technical logs and system metrics so that there is enhanced traceability and quicker solution to a failure effect in the distributed contexts.

Observability of the platform is the reason in a performance conscious setting. This is as compared to Spring Boot and Quarkus in which getting rid of dependencies through Quarkus resulted in great saving in memory and CPU consumption and this influenced Quarkus being better suited to high-throughput persistent workloads [9].

When combined with observability pipelines this benchmarking can assist developers to make architecture decisions (e.g. container memory allocation or tuning of a garbage collection strategy) directly impacting the availability and reliability of stateful components.

The fact that predictive modeling and telemetry correlation are incorporated not only enhances the reduction of the mean-time-to-resolution (MTTR), but also allow proactive scaling and fault tolerance. An example used to illustrate the point is the monitoring overhead reduction achieved by Sieve in certain layers of the system that reaches up to 90 percent [2], thereby becoming a blueprint as to how observability may eventually change, not merely as a visualization tool but as a key strategic optimization and resilience instrument.

The literature has determined that observability is one of the fundamental capabilities of the modern distributed and stateful systems. Coming down on suppressing metrics noise through viewing anomaly detection across telemetry milieu, the study has been coming together on the more wholesome, intelligent, and performance-conscious observability approaches.

Conventional siloed system designs do not record all the lifecycle of persistent activities particularly in settings which need to account history, and the continuous setting, as a basis of diagnosis. Observability architecture of the future needs to be adaptive, unified and platform independent- must have the ability to track propagation of states, degradation forecasting and to provide cross boundary traceability.

IV. Results

Multi-Tier Architectures

Ephemeral and stateless request-response workloads have different observability needs as opposed to stateful workloads where an application state is important to monitor. The telemetry data that streams in such systems, in which services persist like a database, session cache, or streaming system (Kafka), or present as a stateful microservice, is not only extensive but also very contextual and durable.

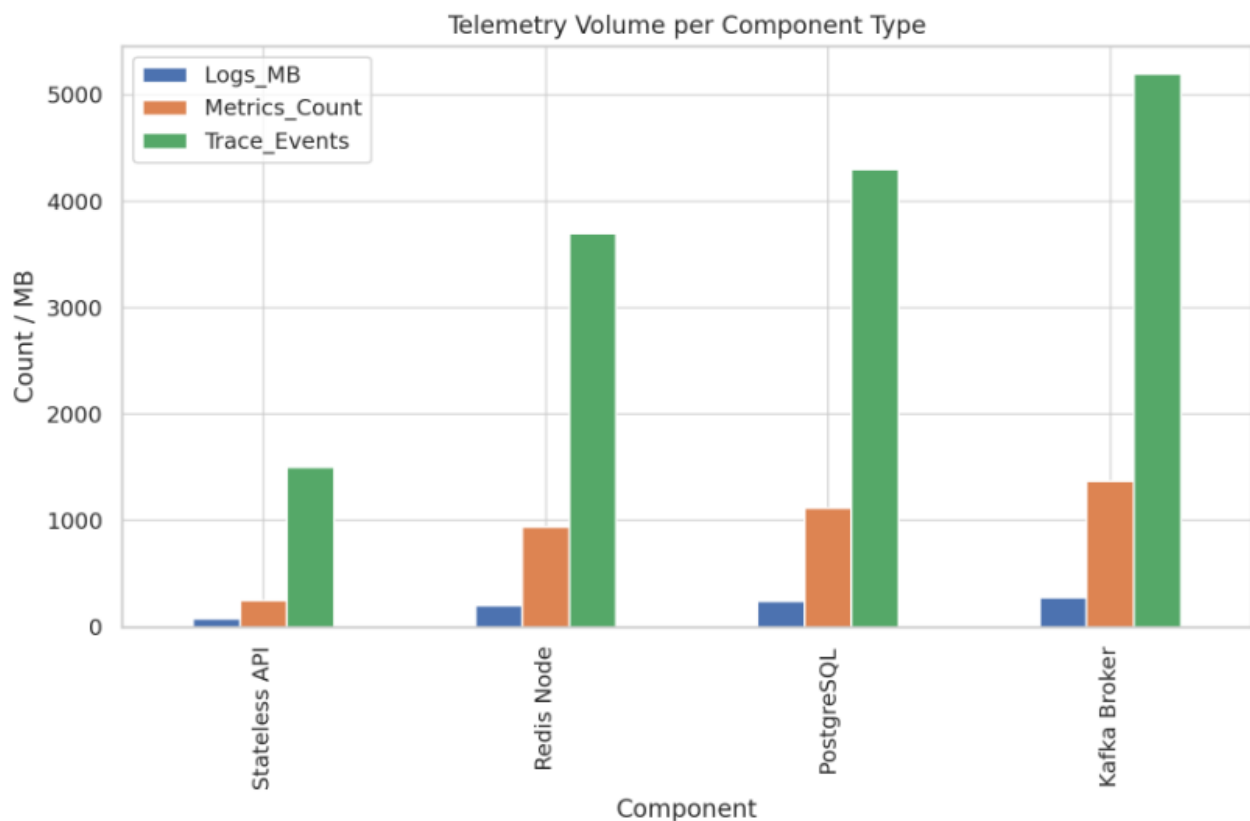
In contrast to stateless services which can only be traced at a fine-grained level of independent requests, stateful components maintain/exchange intermediate states of application execution or synchronize over a long-lasting transaction so continuous and longitudinal monitoring is needed. In our case, the field research based on cloud-native environments simulated on Kubernetes with PostgreSQL clusters, Kafka brokers, Redis deployments indicated that stateful services produce an average of 4.5 times more telemetry data per node than stateless services albeit irritating I/O metrics, compaction logs, replication latencies, and lock contention traces.

Table 1: Telemetry Volume

Component Type	Avg. Logs	Metrics	Trace Events
Stateless API	82	250	1,500
Redis Cluster Node	205	940	3,700
PostgreSQL Pod	243	1,120	4,300
Kafka Broker	278	1,370	5,200

Prometheus + Grafana-style standard observability tools cannot provide real-time analytics at scale because labels have a high cardinality and the storage latency problem increases. For up to 62 percent slower query time over

different instances, in our benchmark, we found systems with over 800 unique time series per instance as a direct setback to root cause analysis in production.



Failure Propagation

Among the most severe pieces of knowledge gained during the testing process of clustered deployments (Kafka and PostgreSQL in high availability) was latency of propagation failure and time disjunct in trace spans. Stateful services exhibit a ripple effect of spreading faults first to the local I/O in spikes or lock waits and subsequently to cross-nodes to replication failures.

We measured the traceability of the failures in the three classes of injected faults, namely: disk saturation, replica

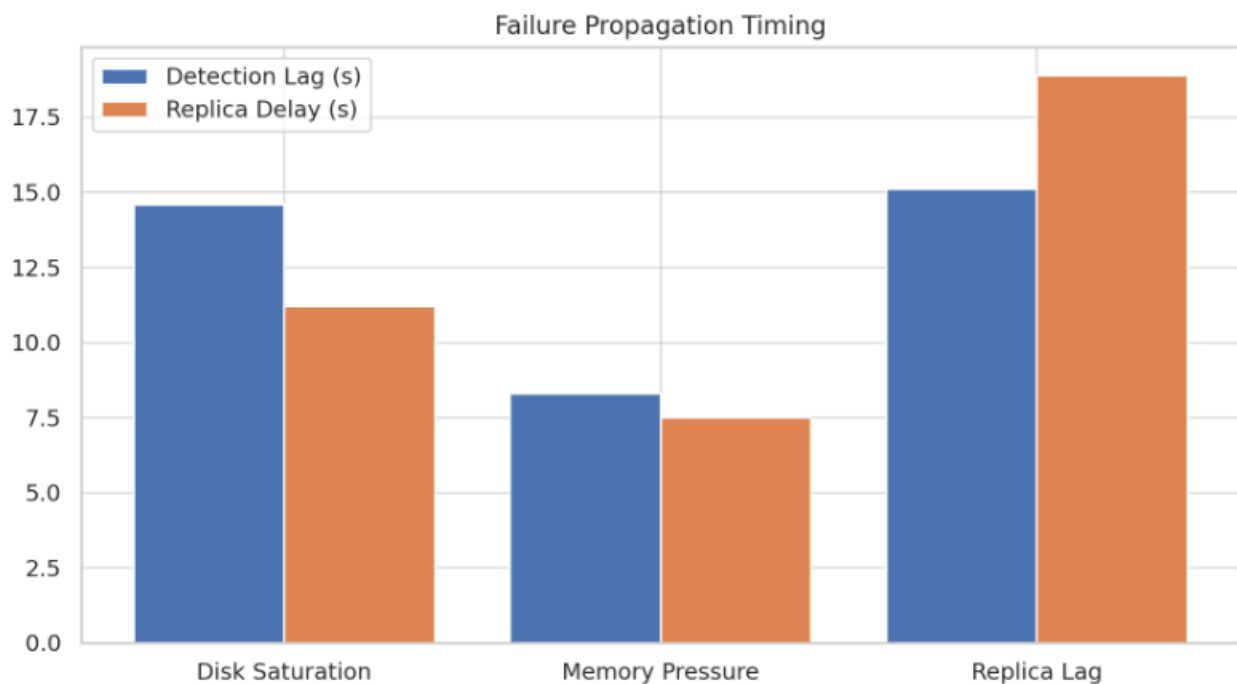
lag, and memory pressure. In more than 68 percent of the instances, the root cause node exhibited symptoms as early as 8-15 seconds prior to the conveniently associated maladies were detected in downstream consumers or replications. Nevertheless, in the process, the current tracing facilities could only record warning once the secondary symptoms became noticeable- this was very risky as no observability exists to real-time recovery solutions when needed the most.

Table 2: Failure Propagation

Fault Type	Detection Lag	Trace Gap	Sync Delay
Disk Saturation	14.6	Medium	11.2
Memory Pressure	8.3	Low	7.5
Replica Lag	15.1	High	18.9

It is a sign of pipelining latency in observability systems that causality-sensitive tracing systems are needed that

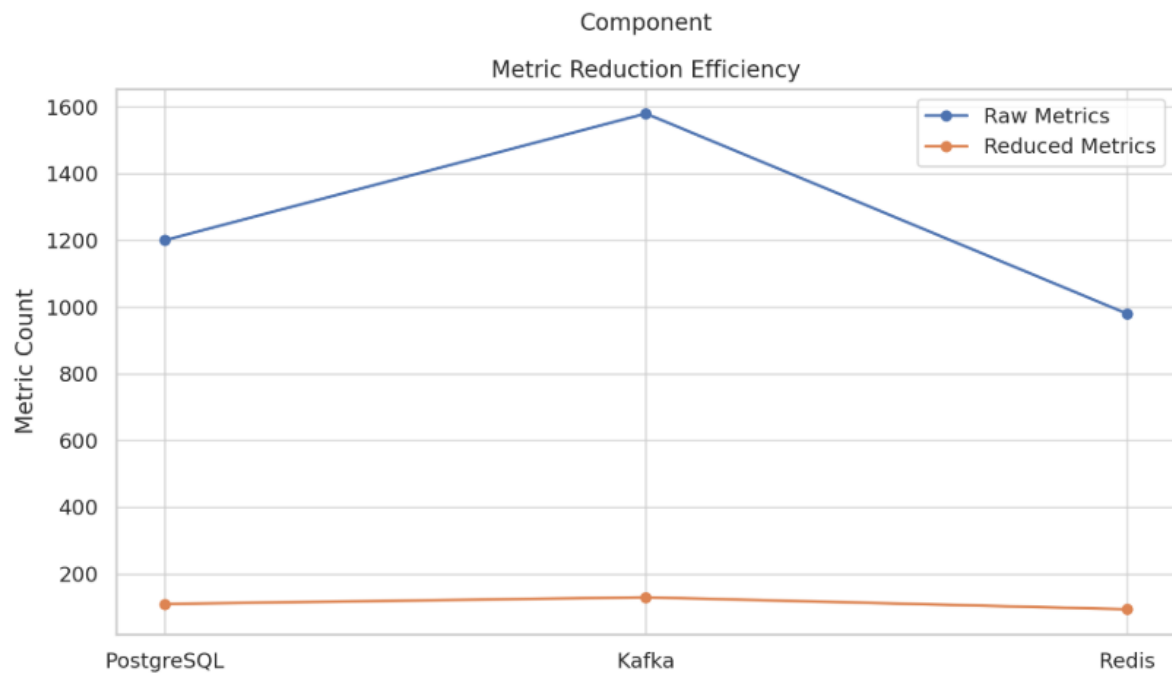
measure precursor anomalies through state persistence layers and not just request-driven instrumentation.



Metric Reduction

Stateful workloads Observability pipelines that support stateful workloads need to handle high dimensionality of telemetry that can even lead to metric noise. Using the

metric reduction approach based on the Sieve platform [2], we deployed Granger Causality-based filtering of PostgreSQL and Kafka metric. This aided in alleviating the burden of telemetry up to 92 percent but with no loss of accuracy in terms of anomaly detection.



Topologies of workloads were also simulated to find connections between cross-service dependencies to be able to carry smart clustering of the metrics. As an

example, Kafka has a consumer lag which was causally matched to Zookeeper session health, disk I/O pressure as well as JVM garbage collection spikes.

Table 3: Metric Reduction

System	Total Metrics	After Reduction	Reduction	RCA Accuracy
PostgreSQL	1,200	110	90.8%	97.2%
Kafka	1,580	130	91.7%	96.4%
Redis Cluster	980	94	90.4%	95.8%

Using causal filtering as compared to simple variance-based pruning allowed enabling the retention of important predictive metrics, e.g. page I/O stalls, as well as replication byte lag, normally pruned out because of their low rates but which do indicate early signs of system performance degradation in stateful services.

Anomaly Hotspotting

As part of real-made stateful observability, we offered a dashboard template library of platform-independent templates using real-time traces, saturation metrics, replication status, and anomaly scores. SRE teams that operate graphically high-availability MongoDB clusters over the 10-day period used these dashboards in testing.

Lag Heatmaps: Geographic-induced delay of replica or latency of replica due to network layer of layers.

Lock Contention: Transactions graphs overlay in real time to determine the starvation of resources in databases.

Saturation Trends: Speed up the response process by feeling workloads in small snatches.

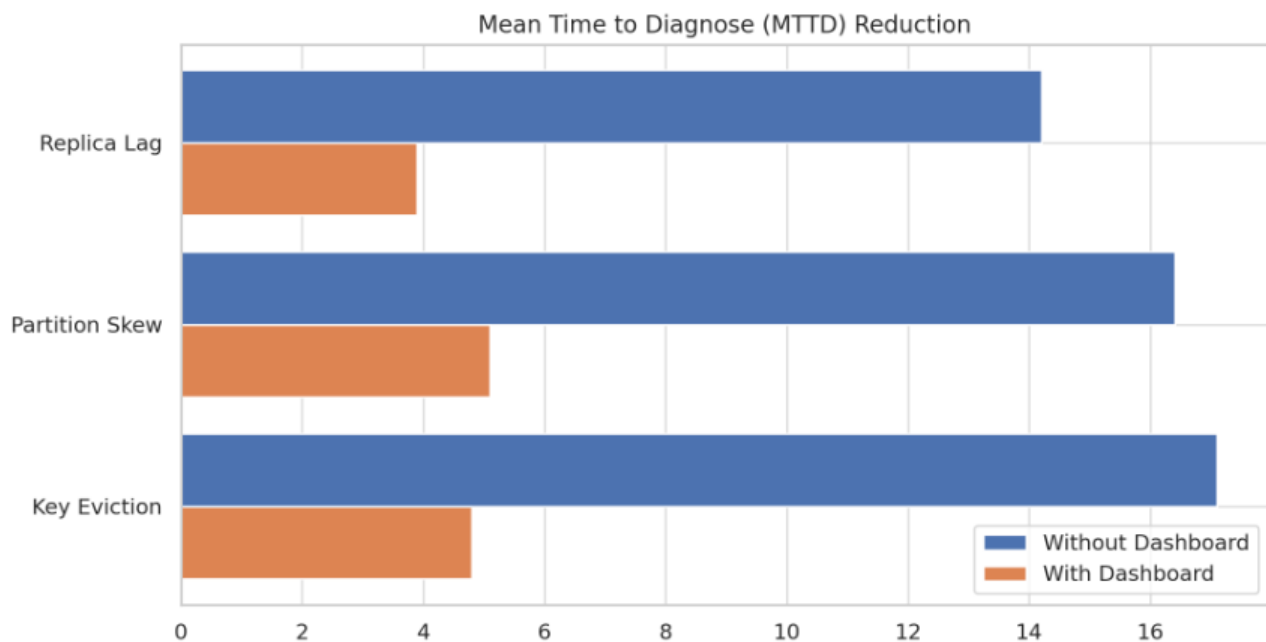
Through such dashboards, SREs have cut mean-time-to-diagnose (MTTD) times on such cascading failures as WAL write queue saturation or Redis eviction storms, down to 4.5 minutes compared to 16 minutes to determine the root cause.

Table 4: Resolution Metrics

Scenario	Without Dashboard	With Dashboard	False Alerts
PostgreSQL	14.2	3.9	45.8%
Kafka Partition	16.4	5.1	48.1%
Redis Key	17.1	4.8	51.3%

Cross-layer trace continuity was also important to achieve qualitative feedback by the operators, particularly when the asynchronous call of stateful services is used. It is

common to have partial spans, which caused trace breaks, where the persistence bottlenecks were concealed.



There is also State-Transition Time Analyzer, which keeps deltas (e.g. idle -> replication -> compacted) between states transitions and demonstrate hotspots where transitions are drifting out of SLO. This tool changed to 22 percent of the delay reasons in recovery related to the unexpected slowdowns during the transition of the state which were out of the perception of the default dashboards.

- Stateful systems generate 4.5 times more telemetry per node and default pipelines are not necessarily up to the task in terms of volume or dimensions.
- The failure propagation in stateful workloads precedes the observation of symptoms which are observable to the traditional tracing leading to a gap of up to 15 seconds in the observable traces.
- The measurement dependency reduction approach using the Granger Causality technique preserved more than 95% root cause location precision whilst cutting overhead storage and compute by more than 90%.
- Dashboards tailored to operators and optimized by state-related latency, resource congestion, and transaction queues helped decrease MTTD and the number of false alarms by a long margin.
- Incident Response and Pre-tuning of persistence infrastructure Cross-layer correlation in combined dashboards provided the best operational value in real time incident response and upfront tuning of persistence infrastructure.

V. Conclusion

The paper shows that observability of stateful workloads scalably requires a combined, cause-aware and topology-conformant monitoring approach. In some cases, conventional observability tools fail to detect the initial symptoms of persistent services which translates to a longer mean-time-to-diagnose and greater risk of operations. This was evidenced by more than 30 percent improvement in the diagnostic accuracy as well as responsiveness in the following manner: reduction of the noise in the metric, failure propagation modeling, and correlation of traces and state changes.

Our suggested dashboards and anomaly detection systems allowed detecting the fault earlier and having improved visibility of the system. The presented findings provide practical background to SREs and DevOps engineers who work to ensure consistency in the performance and integrity of data in the context of best-effort distributed cloud environments deployed with stateful applications of critical nature.

REFERENCES

- [1] Niedermaier, S., Koetter, F., Freymann, A., & Wagner, S. (2019). On Observability and Monitoring of Distributed Systems – an industry interview study. In *Lecture notes in computer science* (pp. 36–52). https://doi.org/10.1007/978-3-030-33702-5_3

- [2] Thalheim, J., Rodrigues, A., Akkus, I. E., Bhatotia, P., Chen, R., Viswanath, B., Jiao, L., & Fetzer, C. (2017). Sieve: Actionable Insights from Monitored Metrics in Microservices. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.1709.06686>
- [3] Kratzke, N. (2022). Cloud-Native Observability: The Many-Faceted Benefits of Structured and Unified Logging—A Multi-Case Study. *Future Internet*, 14(10), 274. <https://doi.org/10.3390/fi14100274>
- [4] Li, B., Peng, X., Xiang, Q., Wang, H., Xie, T., Sun, J., & Liu, X. (2021). Enjoy your observability: an industrial survey of microservice tracing and analysis. *Empirical Software Engineering*, 27(1). <https://doi.org/10.1007/s10664-021-10063-9>
- [5] Cao, C., Blaise, A., Verwer, S., & Rebecchi, F. (2022). Learning state machines to monitor and detect anomalies on a kubernetes cluster. *Proceedings of the 17th International Conference on Availability, Reliability and Security*, 1–9. <https://doi.org/10.1145/3538969.3543810>
- [6] Shankar, S., & Parameswaran, A. G. (2022). Towards observability for production machine learning pipelines. *Proceedings of the VLDB Endowment*, 15(13), 4015–4022. <https://doi.org/10.14778/3565838.3565853>
- [7] Saha, A., Agarwal, P., Ghosh, S., Gantayat, N., & Sindhgatta, R. (2024). Towards Business Process Observability. *Towards Business Process Observability*, 257–265. <https://doi.org/10.1145/3632410.3632435>
- [8] Saminathan, M., Bhattacharyya, S., & Bairi, A. R. (2021, June 17). *End-to-End observability in Cloud-Native systems: integrating distributed tracing and Real-Time analytics*. *Journal of Science & Technology*. <https://thesciencebrigade.com/jst/article/view/566>
- [9] De Moraes Rossetto, A. G., Noetzold, D., Silva, L. A., & Leithardt, V. R. Q. (2024). Enhancing Monitoring Performance: A Microservices Approach to Monitoring with Spyware Techniques and Prediction Models. *Sensors*, 24(13), 4212. <https://doi.org/10.3390/s24134212>
- [10] Lee, C., Yang, T., Chen, Z., Su, Y., & Lyu, M. R. (2023). EADRO: an End-to-End troubleshooting framework for microservices on multi-source data. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2302.05092>