

## Cloud Observability: A Framework for Modern Cloud Architecture

Siva Kumar Mamillapalli\* and Ramya Devi Jeganathan\*

USA

### ABSTRACT

This research investigates the effectiveness of cloud observability tools in improving the performance and reliability of cloud applications. A key challenge is the limited understanding of how these tools affect system metrics and user experience under diverse operating conditions. This study addresses this gap by collecting and analyzing quantitative performance data, user interaction logs, and incident reports from multiple cloud environments. The findings demonstrate that cloud observability tools not only enhance application performance metrics but also improve user satisfaction by reducing latency and downtime. More broadly, this research suggests that improved cloud observability can lead to greater reliability in digital services across various sectors, increasing user trust and encouraging cloud technology adoption. By demonstrating the tangible benefits of these tools, this study provides valuable recommendations for organizations seeking to leverage cloud solutions for enhanced service delivery, contributing to better outcomes and greater operational resilience in a digital world.

This paper presents a technology reference framework for implementing observability. The framework outlines the key functional technology areas that need to be considered for creating actionable feedback loops that use telemetry data for data center and cloud deployments. The technology framework can be used as guiding reference for evaluating and optimizing existing monitoring architectures and implementing new observability solutions.

### \*Corresponding author

Siva Kumar Mamillapalli and Ramya Devi Jeganathan, USA.

**Received:** November 02, 2022; **Accepted:** November 09, 2022; **Published:** November 16, 2022

**Keywords:** Cloud, Observability, Cloud Native Applications, Monitoring, Telemetry, Micro Services Architecture, Event Correlation

### Introduction

In recent years, the rapid evolution of cloud computing has fundamentally transformed how organizations manage applications, data, and IT resources, creating an urgent need for effective observability solutions. With the growing trend of organizations migrating to cloud environments, the complexity of managing distributed systems not only increases but also compounds the challenges in monitoring performance, diagnosing issues, and ensuring reliable service delivery. This situation highlights a critical research problem: despite the paramount significance of cloud observability, many organizations encounter substantial hurdles when trying to implement comprehensive monitoring solutions that can adequately capture real-time insights under varying operational conditions. This research seeks to systematically address this problem by investigating not only the effectiveness of various cloud observability tools but also evaluating their impact on system performance and user experience, while simultaneously identifying best practices for their implementation. The primary objectives of this dissertation encompass a critical evaluation of the capabilities and limitations of widely used observability tools such as Grafana, Prometheus, as well as an exploration of how these tools can enhance incident response, optimize application performance, and foster seamless user experiences. Moreover,

this research aims to undertake a thorough analysis of specific case studies across diverse sectors, where system reliability and real-time data accessibility are of utmost importance. The significance of this exploration is profound, as it has the potential to provide actionable insights for both academics and practitioners; the findings may not only influence future research directions but also inform practical applications in cloud observability. By critically examining effective monitoring strategies, this dissertation aspires to guide organizations in strategically employing observability as a fundamental principle of their cloud infrastructure management. This approach is expected to lead to enhanced operational efficiency and improved user satisfaction. Furthermore, considering the increasing reliance on cloud services, there is a pressing need for a more nuanced understanding of how observability tools can effectively reduce downtime and bolster service reliability, which further amplifies the urgency of this research. The relevance of these discussions is further emphasized by the deployment of multifaceted observability frameworks that integrate various data sources and monitoring protocols to cultivate a resilient IT ecosystem. This is evidenced in system architecture diagrams and data management practices illustrated in several cloud observability examples. By aligning observability efforts strategically with organizational business objectives, companies can not only navigate the complexities of cloud environments with greater assurance but also drive continuous improvement in service delivery metrics, thus solidifying their competitive advantage.

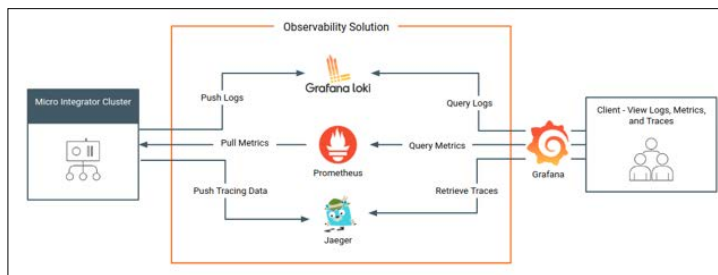
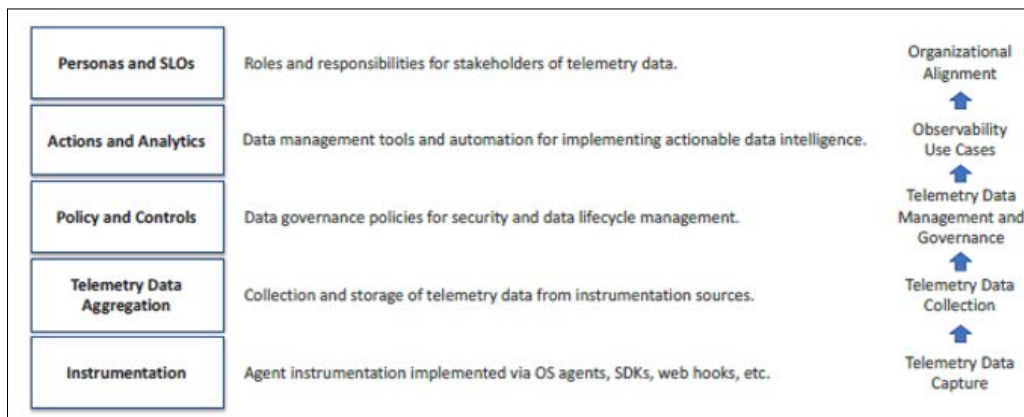


Image 1: Observability Solution for Micro Services Architecture

**Observability Drivers**

The following image depicts cloud observability framework for assessing and implementing observability solutions. The framework highlights the essential functional areas to consider when establishing actionable feedback loops that utilize telemetry data for both data center and cloud environments. We define telemetry data consumers as personas and frame use cases as Service Level Objectives (SLOs).



Over the past decade, there has been a strong push towards enhancing enterprise IT agility to support the digital transformation of business operations and services. This shift has catalyzed various technological advancements such as infrastructure virtualization and containerization. These innovations have emerged from significant changes in the software development landscape, moving away from monolithic architectures towards more distributed implementations using microservices. This transition enables faster feature deployment and greater flexibility through software-defined environments, facilitating rapid release pipelines that integrate business logic and platform components.

However, alongside these benefits of agility and speed, complexity has increased noticeably. For instance, new orchestration tools like Kubernetes and cloud services simplify the provisioning of IT environments and deployment of application runtimes, yet they also add layers of complexity, particularly in terms of monitoring and management requirements. Given the frequent updates in infrastructure and applications, timely feedback on deployment status and rapid detection of failures are crucial. The traditional approach of using isolated IT monitoring tools is proving insufficient in handling the volume and complexity of telemetry data generated by modern containerized microservices deployments. This has spurred a shift towards "Observability" within the cloud-native movement, emphasizing a comprehensive, system-oriented approach to IT monitoring. Observability integrates various monitoring data types to produce actionable insights, becoming a critical component in cloud-native deployments.

Similar trends are observed across other industries like aviation and automotive, where increasing system complexity necessitates advanced monitoring systems for real-time control and decision-making. In summary, the focus in IT monitoring has evolved from mere data collection towards leveraging heterogeneous telemetry data to derive actionable insights and enable automation, akin to the "driver assist" systems in modern vehicles. The terms observability and monitoring are often used interchangeably but they have different meanings. Observability has its origin in control theory where it refers to the ability to derive the (internal) state of a system based on its external outputs. The goal is to externalize the system state based on sensor data outputs.

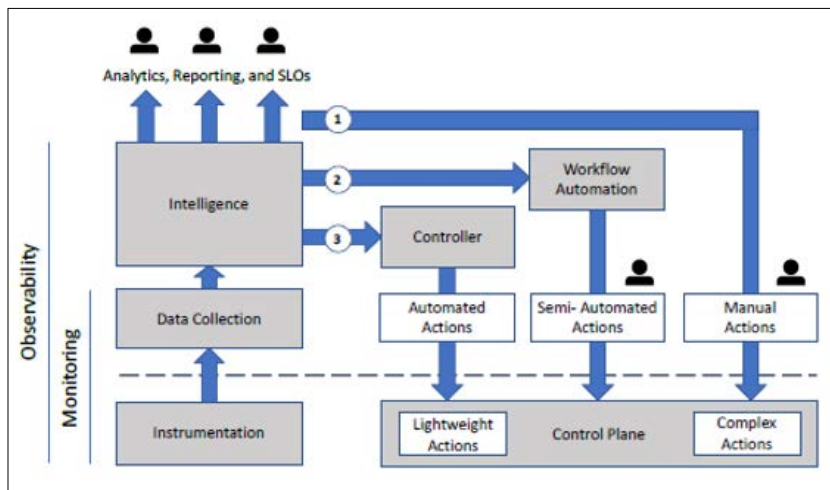


Image 2: Observability Feedback Loops and Monitoring

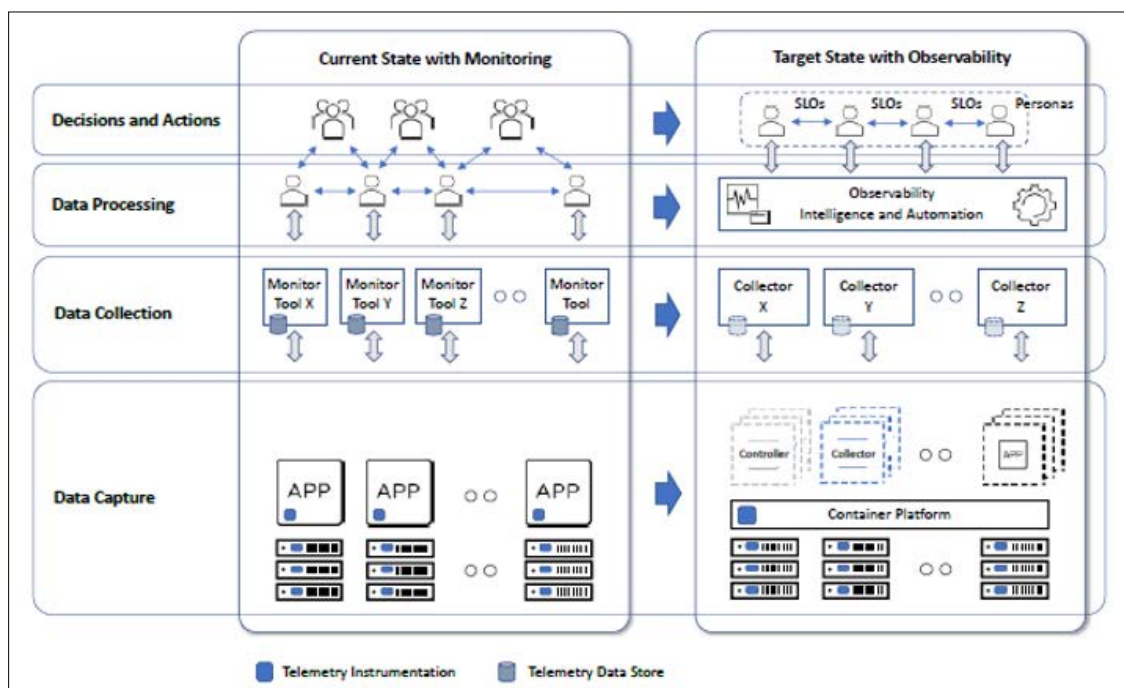


Image 3: Transitioning from Monitoring to Observability

### Cloud Observability Framework

The cloud observability framework serves as a guide for evaluating current monitoring architectures and implementing new observability solutions using monitoring tools, cloud services, or a combination of both. Each function within the framework can be treated as a distinct area of focus for design and implementation decisions. The main components of the cloud observability framework are as follows:

- **Personas and SLOs:** Personas and Service Level Objectives (SLOs) help align telemetry data usage across stakeholders. Personas define roles and responsibilities within the IT organization, while SLOs provide a blueprint for establishing agreed-upon quality levels for IT services and applications that these personas manage. To monitor SLOs, Service Level Indicators (SLIs) are defined, which translate into specific metrics that can be collected and tracked over time.
- **Actions and Analytics:** Telemetry data is used in observability for operational tasks such as troubleshooting, planning, trending analysis, SLO status reporting, and automating platforms and workflows. To make this data actionable, large data sets need to be aggregated and processed, often requiring advanced analytics tools and search capabilities.
- **Data Policy and Controls:** When telemetry data storage is required, data governance policies addressing security and lifecycle management must be established. Security controls ensure that only authorized personas can access the relevant data, often through an enterprise identity management system. Data access policies are especially crucial in multi-tenant environments, and retention policies must be defined to manage the growth of telemetry data, particularly for containerized microservices. Different retention intervals may apply to telemetry data from different environments (e.g., development, testing, production).

- **Telemetry Data Aggregation:** Telemetry data is collected from various instrumentation points. The collection process involves connectivity and security (authentication and encryption), data ingest processing (e.g., field transformations), and storage. Business continuity requirements influence the design of connectivity and storage, ensuring both resiliency and transport continuity, such as dual data collection feeds. Cloud monitoring services also impose additional security requirements for telemetry data sent to and from external endpoints.
- **Instrumentation:** Agents are deployed to capture telemetry data from various layers in the technology stack. With the shift towards cloud-native architectures, telemetry data collection methods now extend beyond traditional OS agents to include containerized metrics collectors, log forwarders, proxy agents, application runtime libraries, web hooks, and service APIs.

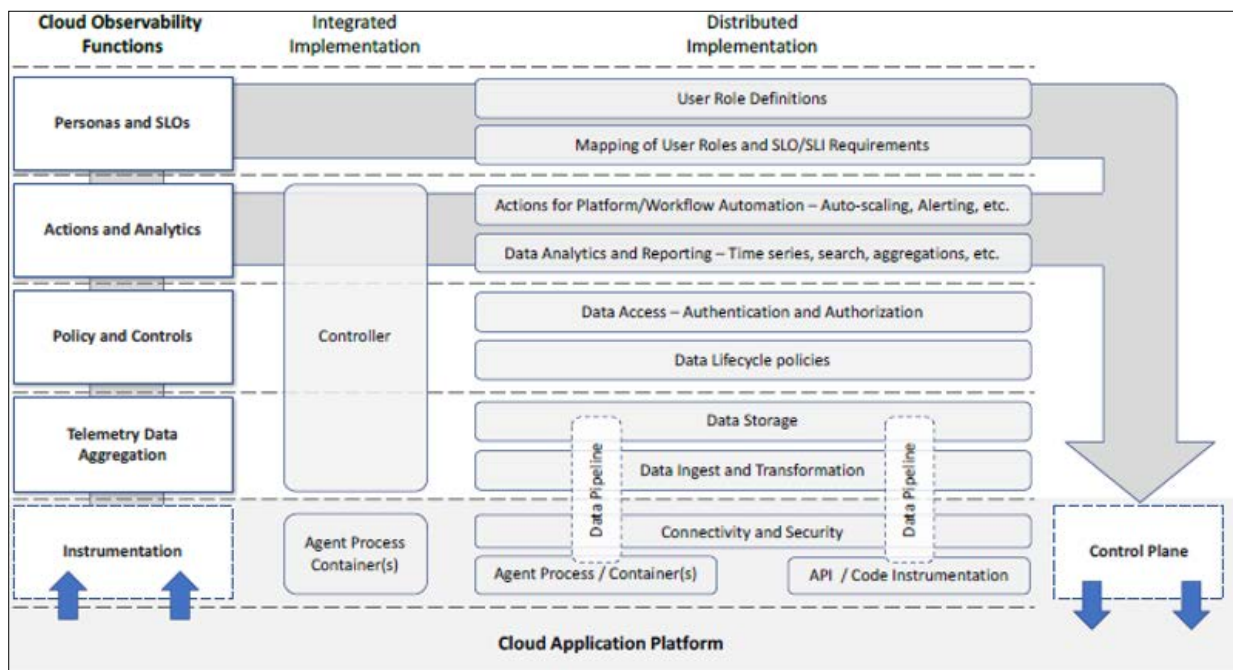


Image 4: Cloud Observability Framework

**Conclusion**

Effective observability in cloud computing environments is paramount to enhancing operational efficiency and reliability, as highlighted throughout this dissertation. The discussion has delved into key themes, emphasizing the integration of observability tools such as telemetry and metrics, while also critically evaluating the roles of artificial intelligence and automation in bolstering real-time monitoring capabilities. The research problem surrounding the challenges of managing complex cloud ecosystems has been thoroughly analyzed, leading to the development of a comprehensive observability framework that prioritizes the effective utilization of data across diverse cloud architectures. This framework demonstrates that organizations which adopt standardized practices and technologies can realize significant performance improvements, which in turn facilitates quicker incident resolutions and more informed decision-making. The implications of these findings extend beyond theoretical concerns; they highlight the essential need for interdisciplinary approaches that effectively integrate observability into existing IT frameworks, thereby establishing a robust foundation for best practices in cloud operations management. From a practical standpoint, this research serves as a valuable guideline for organizations aiming to leverage cloud observability proficiently, thereby enhancing their capabilities in resource utilization and incident response. Furthermore, recommendations for future work suggest extending this framework's application across various sectors and environments, ensuring that continuous improvement is inherently built into observability methodologies. It is also critical to critically assess how adaptable observability strategies can be in response to emerging technologies and evolving business landscapes,

thereby encouraging additional exploration into the intersection of AI and observability tools to drive efficiency. These findings advocate for continuous research into the intricacies of data governance in observability, particularly as organizations strive to maintain compliance and bolster security within increasingly complex cloud infrastructures. Moreover, employing advanced data visualization techniques within observability frameworks, as discussed in the extracted knowledge, could significantly enhance user interaction and provide deeper insights into system performance. Consequently, this dissertation not only contributes to the existing body of knowledge regarding cloud observability but also positions itself as an essential resource for future scholarly discourse and practical application in IT operations, urging readers to consider the broader implications of their findings.

**References**

1. N Kratzke, PC Quint (2017) "Understanding cloud-native applications after 10 years of cloud computing—A systematic mapping study", J. Syst. Softw 126: 1-16.
2. The Twelve Factor App (2021) Available: <https://12factor.net>
3. Senapathi, J Buchan, H Osman (2018) "DevOps capabilities practices and challenges: Insights from a case study", Proc. 22nd Int. Conf. Eval. Assessment Softw. Eng 57-67.
4. P Duvall, S Matyas, A Glover (2007) Continuous Integration: Improving Software Quality and Reducing Risk, Reading, MA, USA: Addison-Wesley Professional.
5. J Humble, D Farley (2010) Continuous Delivery: Reliable Software Releases Through Build Test and Deployment Automation, Reading, MA, USA: Addison-Wesley Professional.

6. A Balalaie, A Heydarnoori, P Jamshidi (2016) "Microservices architecture enables DevOps: Migration to a cloud-native architecture", IEEE Softw 33: 42-52.
7. Cloud Native Computing Foundation, Jan. 2021. Available: <https://www.cncf.io>.
8. J Kosińska, K Zieliński (2020) "Autonomic management framework for cloud-native applications", J. Grid Comput 18: 779-796.
9. N Marie-Magdelaine (2021) "Observability and resources managements in cloud-native environnements" Available: <https://theses.hal.science/tel-03486157>.
10. B Scholl, T Swanson, P Jausovec (2019) Cloud Native: Using Containers Functions and Data to Build Next-generation Applications, Sebastopol, CA, USA:O'Reilly Media.

**Copyright:** ©2022 Siva Kumar Mamillapalli, Ramya Devi Jeganathan. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.