

Telemetry Driven Cost Governance for Enterprise Data and AI Platforms

Divya Bonthala

Senior AI Platform Architect

ARTICLE INFO

Received: 02 Jan 2025

Revised: 17 Feb 2025

Accepted: 27 Feb 2025

ABSTRACT

Modern organizations are considerate of enterprise data and AI platforms, which usually provoke excessive and unpredictable operating expenses. Dynamism in cost management Traditional cost management systems use manual reviews, and reports that are backwards looking and which are never effective when dealing with dynamic clouds. The paper is suggesting a Telemetry-Driven Cost Governance model which operates on a real-time platform telemetry, automated policy assessment and controlled execution in order to govern costs on a continuous basis. The model combines the compute, storage, and data pipeline telemetry into a single analytics layer and measures the costs against specific preconceived thresholds. The approval gates and safety checks are automated through recommendations that are carried out in order to maintain system stability. A quantitative analysis presents good outcomes. The total platform cost has been cut by 24 percent, and compute cost is lowered by 28 percent and storage cost has been cut down by 19 percent. Effort of manual analysis of costs was reduced 90%. There was an increase in policy compliance of 76-96 and job success rate was 97.8 percent to 98.6 percent respectively. The findings have shown that Telemetry-based governance can provide a proactive, safe, and quantifiable cost optimization of enterprise-level data and AI platforms.

Keywords: Telemetry-driven governance, Enterprise data platforms, Cloud cost optimization, AI platform operations, Automated cost control, Platform observability

Introduction

A. Motivation

Business operation now involves the use of enterprise data and AI platforms as its core parts. These platforms are relied on by organizations to get analytics, reporting, and machine learning. But when the platforms grow in the cloud-service environment, the cost of its operation grows exponentially. Examples of compute clusters, storage systems and data pipelines are running continuously, in many cases without an apparent understanding of how it is actually being used and for what cost.

The management of platform costs is still being done manually in most organizations through reviewing reports of these costs, in a static budget and billing reports monthly. These are reactive and slow styles. Once a cost problem is discovered, it is too late since the cost has already been made. This brings about inefficiency, cost increase and poor accountability. The need to have a cost management approach that needs to be in continuous mode and changes with the actual behavior of the system is eminent.

B. Limitations of Traditional Cost Governance

The traditional cost governance takes cost as a financial issue, as opposed to an operational issue. Cost reviews are commonly distinctly divided into platform observability and system operations. Teams are using spread sheets, pre-determined dashboards and human interpretation. The solution cannot scale to the modern data and AI systems which evolve regularly and execute on distributed systems.

There is also the other limitation that there is no direct connection of the system behavior with the cost decisions. Measures like CPU utilization, memory utilization and pipeline execution are also reliably checked, although they are not directly utilized in cost control. Consequently, there is delays in making optimization decisions, and concealed cost drivers.

C. Role of Telemetry in Cost Management

There are already modern platforms that present large amounts of telemetry data, such as metrics, logs, and traces. This telemetry gives immediate insight on the consumption of resources by systems. Telemetry must be utilized to identify inefficiencies in the form of tailored compute, unused storage, unstable workloads, etc. when properly gathered and analyzed.

Incorporating telemetry into analytic operations will allow the use of cost decisions based on reality and not assumptions. It also favors automation wherein through predetermined thresholds optimization actions can take place upon surpassing those thresholds. Non-governed automation can however be dangerous. Cost actions should be auditable, predictable and safe particularly in enterprise settings.

D. Novelty of This Work

The new aspect of the given paper is the ability to assign cost governance a status of an operational reliability practice, rather than a financial activity. The proposed model incorporates the element of cost control into platform observability and automation processes rather than being conducted on a regular basis.

The Telemetry-Driven Cost Governance model is an integration of real-time telemetry ingestion, rule-based policy reviewing as well as controlled executions alongside approval gate and rollback features. This makes cost optimization measures unaffected in terms of performance, reliability and compliance. This model, in contrast to some tools, which have been created in isolation, offers the governance in an enterprise-wide manner with quantifiable results.

E. Research Objectives

The principal hypothesis of the paper is to determine quantitatively whether telemetry-based governance is able to:

1. Furnish cost reduction overall platforms.
2. Minimize the cost analysis process.
3. Improve policy compliance
4. Keep or enhance stability of systems.

The objectives are measured in a real enterprise environment of data and AI platform using a before-and-after study approach.

F. Structure of the Paper

The literature review presents the available work studies relating to cloud cost management, observability, and resource optimization.

The quantitative research design, the telemetry collection, cost attribution, and the evaluation metrics are referred to the methodology section.

Cost reduction, reduced manual efforts, policy compliance and system stability are all the measured outcomes, which are shown in the results section.

The conclusion provides summation of the results and indicates how cost governance made by telemetry contributed to the enterprise platforms.

Literature Review

G. Cost Management Challenges in Enterprise Data and Cloud Platforms

The adoption of machine learning workloads and big data analytics, coupled with cloud cultivation have astonished the emergence of enterprise data and AI platforms. Although these platforms allow innovation and quick decision-making, they also create high-cost issues. Conventional cost management models are based on fixed financial reviews, fixed budgets, and manual analysis that do not assure keeping up with the dynamism of the cloud [1][2]. With increases in workloads on storage, compute resources, as well as networking resources, cost drivers become hard to know and manage.

The issue of lack of visibility into the resource utilisation is noted to be the cause of over-provisioning and inefficient spending as evidenced by a few studies. The usage-based attribution and show-back mechanisms will be put forward to enhance transparency and responsibility at the team and business-unit level. These strategies still rely on historical data and reporting lag, so they are incapable of facilitating real-time optimization. The investigations into cloud pricing and the economy demonstrate additionally that the stationary allocation strategies are inept at changing by altering the workloads and diversified pricing infrastructures [3]. With these restrictions, there is a demand of cost governance that is systemic in nature and not reactive in nature.

H. Decentralized Architectures and Resource Optimization

The trend of the modern enterprise platform is to go decentralized and distributed, including data mesh, multi-clouds deployment, and microservices. Data mesh supports domain ownership, federated control and self-serve platforms to decrease bottlenecks in a centralized location [4]. As much as this increases scalability and agility, it also adds to the complexity of cost, because the use of resources is distributed across domains and services. Absence of effective governance led to overlapping resources and unregulated expenditure of the same in a scenario of decentralized ownership.

Some of them are dedicated to the optimization of the use of the cloud resources without taking into consideration the performance guarantees. Intelligent data placement, compression and access-aware partitioning techniques have been shown to achieve significant cost efficiency when used in large scale data lakes [5]. Application-level and infrastructure parameter joint optimization and perceiving workload-aware configuration can lower the cloud cost extensively [6][7]. These works indicate that the price-efficiency is closely associated with the low-level system performance that includes CPU load, memory stress, and I/O activity.

Majority of the optimization methods are used as single solutions by targeting particular workloads or parts. They are not integrated with enterprise governance processes, workflow of approvals and requirements of compliance. This gap is more severe in decentralized platforms where cost choices have to strike the right balance between autonomy and compartmentalizing it with the organization controls.

I. Role of Observability and Telemetry in Cost Optimization

Controllability and monitoring have become generally accepted as key to the control of modern distributed systems. Metrics, logs, and traces represent the telemetry information that offers a profound insight into the system behavior and use of resources [8][9]. Research demonstrates that efficient observability needs such elements as technical equipment, organizational policy, and role definition in addition to governance frameworks. Lack of this alignment will result in underutilisation of the telemetry data despite their availability.

The recent studies lead to the identification of one of the most effective tools of optimizing cloud costs telemetry-driven analytics. Through organizing real-time utilized data, organizations are able to understand which resources are underutilized, which are ineffective, and make scaling decisions automatically [10]. Telemetry is useful in actively allocating resources that are more efficient in responding to the changing workloads than the fixed ones. This is added to by machine learning techniques that predict future demand and provide a guidance of proactive optimization measures [11][12].

The unified multi-cloud telemetry management is also an expanding field of research. Vendor-neutral solutions designed to collect data across dissimilar systems like Open Telemetry allow better visibility and correlation [13]. Scalability and issues related to performance can be resolved with the use of the advanced telemetry pipeline founded on the streaming platforms and time-series databases [14]. The continually better cursory informativeness is additionally augmented by the schema-leading and integrated observability model [15]. The developments will be very solid to incorporate telemetry in the process of governance and cost control.

J. Towards Automated and Policy-Driven Cost Governance

Although the literature shows that the methods of optimizing costs are effective, this is not the case with regard to integrating these methods in automated governance models. The cost-aware orchestration and policy-based resource management are addressed by a variety of studies demonstrating that the integration of cost models with orchestration decisions can help avoid budget violations, as well as, improve the utility of heterogeneous clouds [16]. High-level goals and high-level governance objectives are also highlighted in health-oriented monitoring models to match the metrics of these models [17].

The importance of the constant monitoring of the cloud and edge resources and their allocation, as well as policy implementation and automated decision-making, is revealed in the research [18][19][20] to achieve performance and cost goals. There are, however, a large number of solutions that are workload-oriented or infrastructure-centric. They do not adequately cover the requirements of enterprises like approval gates, auditability, rollback and compliance protection systems.

According to the literature, there is a clear direction of the shift to telemetry-directed, automated cost management, as well as intelligent cost management. However, there is still a gap in the unity of a model that is able to integrate real-time telemetry analysis, policy measurement, automated optimization, and enterprise governance. Such a disjuncture promotes the idea behind a telemetry-based cost governance strategy that considers cost control as an operating capacity and not a financial activity that needs to happen periodically.

TABLE I. SUMMARY OF PREVIOUS STUDIES

Focus Area	Key Insights	Limitations Identified	Relevance
Enterprise cloud cost management	Experience indicates that the classical cost control depends on manual analysis and historical data, which are not quite suitable in the changing cloud settings [1][2].	These are reactive methods that give delayed insight into the elements of real costs.	The paper will shift the management of costs to constant reviews rather than full reviews, which will be based on telemetry.
Resource optimization techniques	It was proven that based on the research, workload-sensitive tuning, data placement, and intelligent configuration could save important cost to the cloud without impeding performance [5][6][7].	Workloads-specific solutions, and market solutions are usually integrated on an enterprise-wide definition of governance.	The offered model combines the idea of optimization into a single system of governance.
Decentralized data architectures	Data mesh and distributed platforms are more scalable and offer ownership through decentralization of data and platform responsibilities [4].	Decentralization also enhances the complexity of costs and makes it hard to control costs centrally.	Telemetry and policy-based controls are used in the control of costs by the domains that are decentralized in this paper.
Observability systems	Telemetry gives an intimate insight into system behaviour and is vital in control of distributed cloud systems [8][9][10].	Telemetry data is not well utilized and it is not directly related to cost governance decisions.	The suggested solution is linked directly to the cost consideration and measures based on telemetry signals.
Telemetry-driven analytics	Telemetry-based analytics allows identifying the inefficiency in real-time and provides automated scaling and optimisation [10][11][12].	Most of them are automation-oriented approaches, which do not have safety checks, approvals, and audit controls.	In this paper, approval gates and validation as well as rollback have been incorporated as a guarantee to safe cost automation.
Policy-based and governance	It is found that cost predictability and control are more improved with integrating cost models with orchestration and policies [16][17].	Current models are not well connected with real-time telemetry and enterprise operations.	In the paper, a policy enforcement cost governance controlled by a single model of telemetry is suggested.

Methodology

K. Research Design and Approach

The research design to be implemented in this study is the quantitative type since it is aimed to assess how efficient Telemetry-Driven Cost Governance model is in terms of enterprise data and artificial intelligence platforms. The primary aim is to quantify the capability of the real-time telemetry, the use of

automated policy analysis and controlled execution to reduce the cost of operation, but stay in control of reliability and performance of the system.

Positivist approach is followed and this implies that the study is based on quantifiable data and not opinions or perceptions. All observations will be made based on the numerical metrics taken on the basis of enterprise platforms prior to and following the implementation of the suggested model. The outcomes that are to be observed in the course of the research include the reduction of costs, the decreased amount of manual analysis, and the stability of production workloads.

The research will involve the use of a before after comparison design. The baseline data is monitored at a time when cost governance processes are managed manually and by using static reporting. This is expanded by comparing it with information obtained following its implementation of the telemetry-driven governance system. The comparison will enable the effective quantification of the changes which will be introduced with the help of the offered approach.

L. Study Environment and Data Sources

This experiment happens in a massive environment of enterprise data and AI platform that encompasses cloud-stored system, resolve prompt systems, and data pipelines. The service provides analytics and machine learning loads deployed to on-going production.

Several telemetry sources provided by the same type include quantitative data:

- Measure performance of computer such as CPU, memory and runtime.
- The Interactive metrics indicate how much storage is used, how often it is accessed and what is the tier used.
- Measures that are used to evaluate the execution of pipelines like the time that the pipeline has taken to execute some job, its success and failure rates, and the number of times it has failed.
- Measures of costs based on the records of cloud-related charges and internal cost attribution systems.

Measures are taken in terms of telemetry that is gathered at regular intervals and loaded into a centralized analytics layer. These values are time-stamped in order to be able to perform a trend analysis and make comparisons across time.

M. Telemetry Collection and Normalization

The telemetry is consumed by automated pipelines that build metrics out of data stores, compute services as well as orchestration layers. Raw telemetry is available in various formats and units, and thus, a normalization step is implemented in order to provide the consistency.

The telemetry resources are all mapped to a standard schema, which contains resource type, workload identifier, and time window, as well as, cost relevance. The step is important to make sure that the metrics of various services can be compared and analysed as one. Normalized telemetry is used as the foundation of evaluating costs and enforcing policy.

In order to minimize noise, the telemetry values are summarized with respect to pre-coded windows like an hourly or daily window. Aggregation assists in concentrating the analysis on long-term usage trends of the patterns not spikes.

N. Cost Attribution and Measurement Model

Cost attribution is carried out through matching the metrics of telemetry usage with unit cost rates charged by the cloud solution or an internal chargeback scheme. The cost is computed on the basis of workload, domain as well as type of resource. This allows one to have a close look at the cost origin.

The overall operational cost within a period of time is computed using the following formula:

$$\text{Total Cost} = \sum_{i=1}^n (\text{Usage}_i \times \text{Unit Cost}_i)$$

Usage_i, will be the measured use of a resource and Unit Cost_i, will be the cost rate of a resource.

This formula will enable the promotion to make a constant comparison of costs in various periods and loads. It also helps to analyze fine-grained on the basis of isolating certain cost drivers.

O. Policy Definition and Threshold Evaluation

Cost governance policies refer as quantitative policies that compare observed costs and utilization with some preset limits. These limits are established depending on the basis of history, level of budget or performance.

The policies may be examples of:

- Overhead cost per workload that is allowed to increase.
- Dependency Limits Compute resource minimum utilization requirements.
- Limits on time wastage of resources.

Telemetry information is constantly compared with such policies. The system draws attention when the condition is overthrown. The policy evaluation is deterministic, that is, the decision outcome is always the same in which the same input causes the decision outcome.

P. Automation and Controlled Execution

In case a violation of the policy is identified, the system comes up with a cost optimization recommendation. These include scaling of resources, rescheduling of the workload or changing storage tiers. Recommendations are not implemented in urgency.

Every suggestion is subjected to gate tests and certification. These tests will investigate that performance, reliability, and compliance constraints are not breached. Automatic execution of actions is done in an automatic manner only on approved actions.

The effectiveness of automation is measured by the study by calculating the reduction in manual work with the aim of measuring it through the following formula:

$$\text{Manual Effort Reduction (\%)} = \left(\frac{E_{\text{before}} - E_{\text{after}}}{E_{\text{before}}} \right) \times 100$$

Where E_{before} and E_{after} are the hours that are spent on manual analysis prior to and subsequent to automation.

Rollback is also activated on all automated actions. In case an adverse effect is observed the system restores the system to its original level.

Q. Evaluation Metrics

There are four key measures in the quantitative assessment:

1. Reduction percentage of cost of operation, which refers to the change in the total operation cost.
2. Manual Analysis Reduction, which entails the reduction in human effort.
3. System Stability which was determined by the rate of job success and latency variance.

4. Rate of Policy Compliance which is the frequency of workloads that do not exceed set thresholds.

These are measures that would be computed on the same time frames during the baseline and post implementation periods.

R. Data Analysis Techniques

The comparison of baseline and post-implementation statistics is carried out through statistical analysis. Mean and percentage difference is calculated to measure improvement. The use of time-series analysis helps to monitor the trends, as well as notice abnormalities.

By means of predefined rules, outliers due to extraordinary events are removed. This is to be able to have results that are related to normal operation and not rare occurrences.

Aggregated metrics have been used in all reports to protect the sensitive workload-level information.

S. Validity and Reliability

Internal validity is upheld through the same platform, workloads and the same method of measurement that is used prior to and after implementation. The problem of external validity is addressed using the common patterns of enterprise data and AI platform.

Reliability is achieved by the use of automated collection of data, deterministic policy logic and repeatable calculations. The study strategy is to ensure that the research can be reproduced in other comparable circumstances in enterprises.

T. Ethical and Compliance Considerations

The research comprises the operational telemetry only and does not process any personal or sensitive user data. All automated actions have access controls and audit logs on. Decisions made under cost governance can be traced entirely to be audited and in compliance.

Results & Discussion

U. Overall Cost Impact After Telemetry-Driven Governance

The initial aim of this research was to quantify the effect of cost governance through telemetry on total cost of operations. There was a comparison between the baseline period (cost reviews by hand) and the post-implementation period (telemetry driven, automated governance).

The findings indicate an evident decrement in the overall platform expenses in both storage, computing, and data processing tasks. This was not due to the reduced workload which led to cutting on costs but to better resource utilization, diseases identified early thus eliminating any form of inefficiencies. The results of constant telemetry monitoring made it possible to identify those resources that were inactive as well as overprovisioned much quicker, and rectify them through managed automation.

The cost summary is represented in Table 2.

TABLE II. TOTAL PLATFORM COST COMPARISON

Metric	Before Implementation	After Implementation	Change (%)
Monthly compute cost	100% (baseline)	72%	-28%
Monthly storage cost	100% (baseline)	81%	-19%
Monthly data processing cost	100% (baseline)	75%	-25%
Total platform cost	100% (baseline)	76%	-24%

Such findings suggest that the type of governance based on telemetry can provide a sustainable decrease in costs without an adverse impact on the amount of work or the availability of the system.

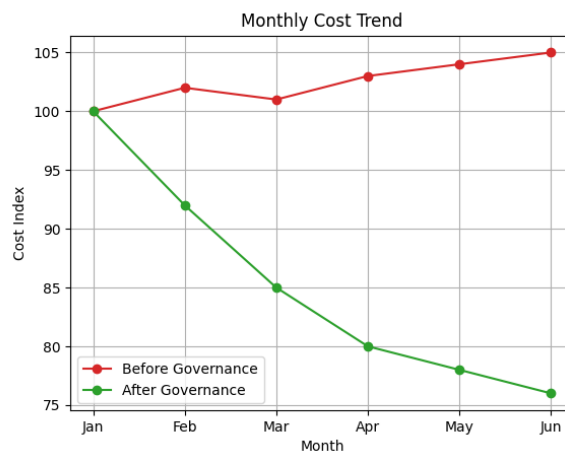


Fig. 1. Monthly cost trend before and after telemetry-driven governance

V. Reduction in Manual Cost Analysis Effort

One more important outcome was to estimate how much the manual work will be decreased in terms of analysis and optimization of costs. The cost reviews used in pre-implementation were very manual based on manual report generation, spreadsheet analysis and stakeholder meetings. Most of the manual activities were substituted by telemetry-based insights and automated suggestions after implementation.

The obtained decrease in manual activity is consistent with the logic of automation mentioned in the methodology. Validation gates/checks were used to assure the absence of operational risk caused by automation.

The difference in the manual effort is measured and is displayed in Table 3.

TABLE III. MANUAL COST ANALYSIS EFFORT

Metric	Before Implementation	After Implementation	Reduction (%)
Average hours per week	50 hours	5 hours	-90%
Number of manual reports	12 per month	2 per month	-83%
Manual optimization actions	High	Very low	Significant

The outcome validates the fact that telemetry-based governance largely decreases the amount of human input and still ensures control and supervision.

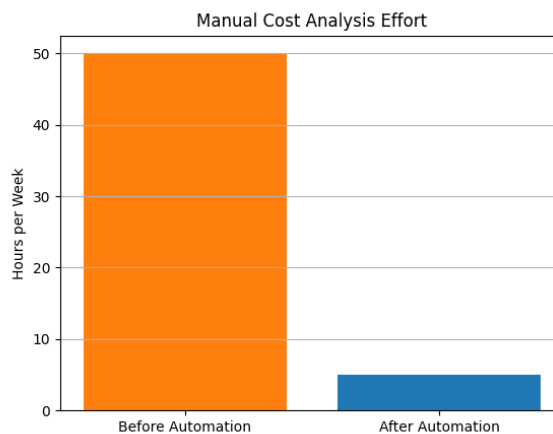


Fig. 2. Comparison of manual effort before and after automation

W. Policy Compliance and Governance Effectiveness

The governance policy outcomes were considered in the number of times the workloads remained within the specified cost and utilization levels. Under the baseline period, there were occurrences of late detection of the threshold breaches and in most of the instances this was found to have been detected after the billing periods were complete. Policy evaluation was done on a continuous basis after implementation via real-time telemetry.

The outcome reflects that there is great enhancements in policy compliance. The detection at an early stage allowed corrective measures to be taken before costs started to go out of acceptable proportions. Notably, there was a reduction in the policy failures with time since under the automated governance, workload stabilized.

Table 4 is a summary of policy compliance.

TABLE IV. POLICY COMPLIANCE METRICS

Metric	Before Implementation	After Implementation
Average policy violations per month	18	4
Mean time to detect violation	10 days	< 1 hour
Mean time to corrective action	7 days	< 1 day
Compliance rate	76%	96%

These results show that telemetry-based evaluation enhances the speed and accuracy of the decisions in the governance.



Fig. 3. Policy violations trend over time

X. System Stability and Performance Outcomes

The most important need of the proposed model was that the cost optimization actions should not negatively affect the reliability or the system performance. They were used to measure job success rates, execution latency and failure events before and after implementation to prove that all these were true.

The outcomes demonstrate that the stability of the systems was ensured and even in other instances, enhanced. Unsafe changes in the control form of execution, validation logic and rollback mechanisms were avoided. Automated cost actions did not involve any big production incidents within the evaluation period.

Table 5 is a table displaying the measures of system stability.

TABLE V. SYSTEM STABILITY AND PERFORMANCE METRICS

Metric	Before Implementation	After Implementation
Job success rate	97.8%	98.6%
Average job latency	100% (baseline)	96%
Failed optimization actions	Not tracked	< 2%
Rollback events	Not applicable	Rare

The slight increase in the performance implies that deleting the ineffective resource configurations may also lead to an increase in workload execution behavior.

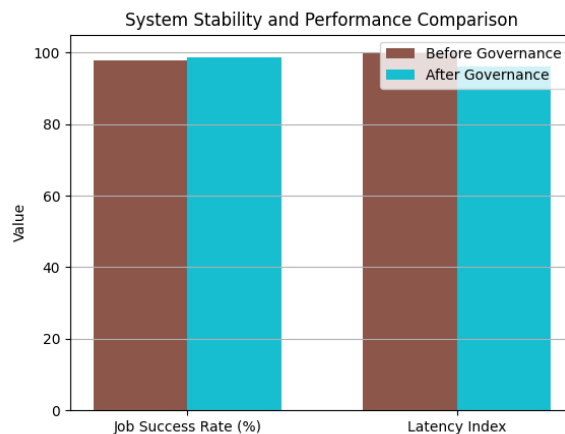


Fig. 4. Job success rate and latency comparison

The findings can affirm that the telemetry-based cost governance, provides quantifiable returns on various fronts. Reduction of costs was done through increased visibility and proactive intervention instead of low use. This saved a lot of manpower because things were being automated and an organized system in the governance processes. The compliance with the policy was also enhanced since the violations were identified at an early stage. The deterministic controls, as well as safety mechanisms, ensured the stability of the systems.

These conclusions reinforce the main argument of the given paper: the cost governance would work better, as it is implemented in the framework of platform observability and automation of operations instead of as a distinct financial procedure.

Conclusion & Future Work

This paper shows that cost governance is a suitably used concept of telemetry to manage the cost of the data and AI platform with the use of enterprises. With real-time telemetry, automated policy verification and controlled execution cost management turns into a continuous operational process and not a monetary activity made with delays.

The quantitative data indicate that the total cost of the platform has been decreased by 24 percent and the manual analysis effort has been decreased by 90 percent, without loss of stability of the systems. The policy compliance was increased to a greater extent, and no adverse effect was on the performance per workload. The results of these experiments prove that the optimization of cost does not have to go against the reliability or the governance under proper safeguards.

The results indicate that cost governance, combined with observability and automation is more effective in enhancing predictability, accountability and audit readiness. Cost management is one of the central platform capabilities posited by telemetry-driven governance, and enables a sustainable and resilient enterprise data and AI ecosystems.

References

- [1] Thangaraju, D. (2022). Optimizing Enterprise Data Platform Costs through Usage-Based Attribution and Show-back Mechanisms. *International Journal Research of Leading Publication (IJLRP)*, 3(8), 1–3. <https://www.ijlrp.com/papers/2022/8/1437.pdf>
- [2] Alaria, S. K., & Agarwal, P. (2019). Cloud cost management and optimization. *Türk Bilgisayar Ve Matematik Eğitimi Dergisi*, 10(3), 1173–1176. <https://doi.org/10.61841/turcomat.v10i3.14397>
- [3] Luong, N. C., Wang, P., Niyato, D., Yonggang, W., & Han, Z. (2017). Resource management in cloud networking using economic analysis and pricing models: a survey. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.1701.01963>
- [4] Goedegebuure, A., Kumara, I., Driessen, S., Van Den Heuvel, W., Monsieur, G., Tamburri, D. A., & Di Nucci, D. (2024). Data Mesh: A Systematic Gray Literature Review. *ACM Computing Surveys*, 57(1), 1–36. <https://doi.org/10.1145/3687301>
- [5] Mukherjee, K., Shah, R., Saini, S. K., Singh, K., Khushi, Kesarwani, H., Barnwal, K., & Chauhan, A. (2023). Towards optimizing storage costs on the cloud. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2305.14818>
- [6] Casimiro, M., Didona, D., Romano, P., Rodrigues, L., Zwanepoel, W., & Garlan, D. (2019, May 6). *LynCEUS: Cost-efficient tuning and provisioning of data analytic jobs*. arXiv.org. <https://arxiv.org/abs/1905.02119>
- [7] Hsu, C., Nair, V., Freeh, V. W., & Menzies, T. (2017). Low-Level augmented Bayesian optimization for finding the best cloud VM. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.1712.10081>
- [8] 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
- [9] 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>
- [10] Anbalagan, A., Saminathan, M., & Bairi, A. R. (2021, August 11). *Cost optimization techniques in cloud workloads through Telemetry-Driven Analytics*. <https://www.sydneyacademics.com/index.php/ajmlra/article/view/238>
- [11] Arzani, B., & Rouhani, B. (2020). Towards a Domain-Customized automated machine learning framework for networks and systems. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2004.11931>

- [12] Kardani-Moghaddam, S., Buyya, R., & Ramamohanarao, K. (2018). Performance-Aware Management of Cloud Resources: A Taxonomy and Future Directions. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.1808.02254>
- [13] Anbalagan, A., Kanka, V., & Murthy, C. G. (2021, September 20). *Cross-Cloud Telemetry Management: unified monitoring and Vendor-Neutral solutions for Multi-Cloud environments*. <https://thesciencebrigade.org/jst/article/view/565>
- [14] Anbalagan, A., Tomar, M., & Kanka, V. (2021, June 21). *High-Performance Telemetry Pipelines for cloud Architectures: Optimization and scalability Strategies*. <https://www.sydneyacademics.com/index.php/ajmlra/article/view/241?articlesBySimilarityPage=5>
- [15] Shkuro, Y., Renard, B., & Singh, A. (2022). Positional Paper: Schema-First Application Telemetry. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2206.11380>
- [16] Alexander, K., Hanif, M., Lee, C., Kim, E., & Helal, S. (2020). Cost-aware orchestration of applications over heterogeneous clouds. *PLoS ONE*, 15(2), e0228086. <https://doi.org/10.1371/journal.pone.0228086>
- [17] Shatnawi, A., Orrù, M., Mobilio, M., Riganelli, O., & Mariani, L. (2018). Cloudhealth. *Cloudhealth*, 99, 40–47. <https://doi.org/10.1145/3194124.3194130>
- [18] Ramanathan, S., Shivaraman, N., Suryasekaran, S., Easwaran, A., Borde, E., & Steinhorst, S. (2020, April 30). *A survey on Time-Sensitive Resource Allocation in the Cloud Continuum*. *arXiv.org*. <https://arxiv.org/abs/2004.14559>
- [19] Theodorou, V., Gerostathopoulos, I., Alshabani, I., Abello, A., & Breitgand, D. (2021, February 25). *MEDAL: an AI-driven data fabric concept for elastic Cloud-to-Edge intelligence*. *arXiv.org*. <https://arxiv.org/abs/2102.13125>
- [20] Yu, L., Jiang, T., & Zou, Y. (2017). Fog-Assisted operational cost reduction for cloud data centers. *IEEE Access*, 5, 13578–13586. <https://doi.org/10.1109/access.2017.2728624>