



COST TRANSPARENCY AND BILLING INTERPRETABILITY CHALLENGES IN AWS ENVIRONMENTS: AN EMPIRICAL STUDY OF BEGINNER CLOUD USAGE

*** Ms. Ruchita Subhash Kamble , **Ms.Tamsil Fatima Firoz Dalvi , ***Mr. Shubham Babu Ponnamm & **** Ms. Vandana Maurya**

*, **, ***, ****, Students, B.K. Birla College (Empowered Autonomous Status), Kalyan.

Abstract:

The use of cloud computing is now the primary infrastructure model available to students, beginning developers and other educational institutions. Amazon Web Services (AWS) created the AWS Free Tier to motivate cloud adoption. Since this Cloud Service (CS) offers limited and free access to core CS, it still presents a problem for new users who struggle to comprehend their AWS billing and the real-time costs involved in starting to use a CS. The purpose of this study is to investigate and evaluate the ideas of billing interpretability and cost visibility for AWS Free Tier. In this scenario, I will recommend a five-stage Cost Interpretability Framework focusing on each user's configuration of services, how those services were used (measured), how those services were viewed displayed on the dashboard, and the various effects of gap-to-gain between each of these factors on a user's behaviour, learning, and experimentation with AWS Free Tier services. This research shows that even if users stay within their AWS Free Tier limits (use limits), there are still interpretability (cognitive) factors related to costs and usage that affect experimentation and behavioural trends. In order to better help students and beginners in embracing and learning how to use Cloud Services, suggestions are made to enhance the dashboard's user interface clarity. **Keywords:** Conceptual Framework, AWS Free Tier, Cost Transparency, Billing Interpretability, New Cloud Users

Copyright © 2026 The Author(s): This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY-NC 4.0) which permits unrestricted use, distribution, and reproduction in any medium for non-commercial Use Provided the Original Author and Source Are Credited.

Introduction:

Cloud computing has grown into an essential part of research environments, educational institutions, and startups. With its Free Tier, which enables limited service usage without instant cost, Amazon Web Services (AWS) has significantly increased beginner adoption. Fortunately, this model eliminates barriers to entry, it additionally introduces issues with billing interpretability and cost transparency. Services like compute, storage, data transfer, and monitoring are billed separately under AWS's granular, usage-based pricing structure. Because of complicated dashboards, delayed billing updates, and technical pricing terminologies, it can be difficult for beginners to understand how actual usage transforms into potential charges.

Beginning users in learning environments frequently report feeling uncertain and concerned of unknown charges in a result of these uncertainty. Users might therefore limit experimentation, which may restrict opportunities for hands-on learning. Additionally, there is a gap in learner-centric analysis because the majority of current research concentrates on enterprise-level cost optimisation. This study identifies interpretability

issues, empirically examines billing accuracy in AWS Free Tier environments, and makes recommendations for improved cost clarity for beginner cloud users.

Problem statement:

Beginner users of the AWS Free Tier frequently struggle to interpret service usage and billing information due to complex pricing structures and limited cost transparency. This uncertainty can lead to perceived financial risk and constrained learning behavior. However, empirical research examining these challenges from a learner-centric perspective remains limited.

Literature Review:

1. Cloud Cost Structures and Pricing Complexity

Cloud computing platforms operate under granular, consumption-based pricing models where users are billed according to measurable resource usage such as compute time, storage allocation, and network transfer. While this model enables scalability and flexibility, it introduces multidimensional billing structures that may increase interpretive complexity. Recent taxonomic research on cloud cost components demonstrates that pricing structures often combine multiple billing dimensions, including service category, usage thresholds, and geographic region (Khan et al., 2024).

Empirical evidence further suggests that cost awareness among developers tends to emerge reactively rather than proactively. Feitosa et al. (2024) observed that cost-related decisions are frequently addressed after deployment or billing review rather than during initial configuration stages. This indicates that pricing transparency alone may not ensure effective cost interpretation, particularly in learning environments.

Although existing studies extensively examine enterprise cost optimization and governance mechanisms, fewer investigations focus on interpretability challenges at the beginner level.

2. AWS Free Tier and Transparency Mechanisms

To support new users, Amazon Web Services provides the AWS Free Tier, offering limited usage across compute, storage, database, and serverless services. The Free Tier includes time-bound and usage-bound needs that differ across services, despite its purpose for reducing financial barriers.

Cost transparency and accountability procedures are highlighted in AWS documentation (Amazon Web Services, Inc., 2025). However, users who are not familiar with limited infrastructure systems may find it challenging to figure out eligibility requirements and service interdependencies. Even inadvertently going over one usage dimension while staying within other Free Tier limitations may result in charges.

For beginners who mainly depend on interface clarity for financial knowledge, this complicated structure may result in a disconnect between documented limits and dashboard-level cost perception.

3. Billing Dashboards and Cost Monitoring Tools

Cloud providers offer monitoring tools to enhance billing visibility. Within the AWS cloud, cost reports are available through the AWS Cost Explorer. These reports serve as a basis for usage metrics, estimated costs, and future predictions of cloud usage. These reports are also essential for supporting governance initiatives and FinOps (Deochake, 2024).



However, since dashboard reporting is typically produced in the form of 'aggregated metrics', these metrics will often be required to be analysed in context. For example, technical units of measure such as vCPU hours, I/O requests, and amount of data transferred may not clearly convey financial meaning to those who are new to AWS cloud usage. Research on serverless computing and cost dynamics has indicated that micro-billing based on usage may lead to non-linear accumulation of charges further complicating understanding of your cloud bills (Hamza et al., 2023).

As a result, dashboards allow users to understand numerical values; however, interpretability of the numerical data presented will depend on the user's level of experience with pricing models and their ability to correctly interpret what they are seeing.

4. Research Gap

Existing literature largely focuses on enterprise-level cost optimization, governance frameworks, and automated FinOps solutions (Deochake, 2024; Khan et al., 2024). Empirical work has examined cost awareness behaviors in development environments (Feitosa et al., 2024) and cost accumulation patterns in serverless systems (Hamza et al., 2023).

However, limited research has specifically investigated cost interpretability challenges experienced by beginner users operating within AWS Free Tier environments. The alignment between documented Free Tier limits and real-time dashboard representation remains underexplored from a learner-centric perspective.

This study addresses this gap by conducting an observational analysis of cost visibility mechanisms within beginner- oriented AWS Free Tier usage contexts.

Research Objectives:

The primary objective of this study is to examine cost transparency and billing interpretability challenges encountered by beginner users operating within AWS Free Tier environments. Specifically, the study aims to:

- Analyze the structural characteristics of AWS Free Tier billing mechanisms and their visibility within cost monitoring dashboards.
- Evaluate the degree of alignment between documented Free Tier usage limits and real-time dashboard-level cost representation.
- Identify common configuration patterns that may contribute to perceived or actual billing ambiguity among beginner users.
- Examine how cost interpretability influences experimentation behavior and perceived financial risk in learner-centric cloud environments.
- Propose structured recommendations to enhance billing clarity and improve beginner cost awareness within AWS Free Tier contexts.

Research Questions:

To address the above objectives, the study investigates the following research questions:

RQ1: How accurately do AWS Free Tier billing dashboards reflect documented service usage limits for beginner users?

RQ2: What interpretability challenges arise when beginners attempt to correlate service configuration with billing outcomes?

RQ3: Which AWS service categories are most frequently associated with perceived billing ambiguity in Free Tier environments?

RQ4: How does perceived cost uncertainty influence experimentation and learning behavior among novice cloud users?


Proposed Cost Interpretability Framework:

1. Design Principles

The proposed framework is grounded in transparency theory, cognitive load theory, interpretability principles, and financial accountability theory. Transparency theory addresses information asymmetry in cloud billing by emphasizing clear mapping between resource usage and billing components. Cognitive load theory explains that multidimensional pricing structures increase user confusion, particularly for novice users, thereby reducing interpretability. Interpretability in information systems requires traceability, causal clarity, and predictability of cost behavior. Financial accountability theory further highlights the importance of consistent and predictable expenses in maintaining user trust. Accordingly, the framework assumes that increased transparency and reduced cognitive load significantly enhance billing interpretability in AWS environments.

2. Framework Architecture

The proposed framework consists of four interconnected layers. Raw operational interactions like EC2 usage, S3 storage, data transfer, and API calls are collected by the Resource Utilization Layer. The Billing Translation Layer addresses these operations by applying regional multipliers, pricing rules, tier thresholds, and metering units. When this translation logic is hidden, interpretability reduces. The Presentation Layer, that involves the AWS Billing Dashboard, Cost Analysis, and usage reports, presents the processed data. The quality of the visualization affects the understanding of the user. Lastly, cognitive processing, the creation of mental cost models, and decision-making behavior are all captured by the User Interpretation Layer. The framework generates the assumption that predictability, traceability, and transparency all impacts overall interpretability, about cognitive load working as a moderating factor.



1. Framework Architecture

a. Cost Interpretability Metrics

To convert theory into measurable variables, the framework introduces four interpretability dimensions:

- 1. Traceability Score (TS):** Ability to trace a billing charge to a specific resource action. Measured by: % of charges clearly mapped to resource ID & Clarity of service-wise breakdown
- 2. Predictability Score (PS):** Degree to which actual cost aligns with user expectation. Formula concept: $\frac{|\text{Expected Cost} - \text{Actual Cost}|}{\text{Expected Cost}}$ Lower deviation = Higher predictability
- 3. Transparency Index (TI):** Composite measure including:
 - Visibility of pricing rules
 - Clarity of billing components
 - Access to usage data
- 4. Cognitive Load Index (CLI):** Measured using:
 - Number of dashboard steps required
 - Number of pricing variables involved
 - User confusion rating

Higher CLI = Lower interpretability.

Overall Interpretability Index (OII): $OII = \frac{TS + PS + TI}{CLI}$

This forms the quantitative validation base.



b. Theoretical Propositions / Hypotheses

H1: Higher billing dashboard transparency improves cost interpretability for beginner AWS Free Tier users. H2: It becomes simpler to understand billing when charges can be clearly correlated with specific resource actions.

H3: Considered billing clarity improves when expected and actual costs are more closely correlated.

H4: Cost understanding is reduced by the increased cognitive load caused by complicated pricing structures.

H5: Beginner users are more prepared to experiment and perceive financial risk as an outcome of lower cost interpretability.

H6: In comparison to core compute and storage services, some AWS services (such as data transfer and monitoring) result in more billing ambiguity.

Methodology:

With the aim to examine challenges with cost transparency and billing understanding in AWS Free Tier environments, this study uses an empirical experimental research design. To make sure accuracy in service usage and billing tracking an operational AWS Free Tier account had been set up. By applying an approach known as deductive, research proceeds from ideas in theory to the construction of predictions and tests based on evidence. Billing interpretability is the dependent variable, and transparency, accountability, and cost predictability are considered separate factors. Interpretability results being believed to be diminished by cognitive load.

All services were from the AWS Free-Tier and limited to usage according to the standard AWS Free Tier usage policy for EC2, S3, Lambda, RDS and CloudWatch. Data collection for both the AWS Billing and usage reports, as well as structured self-reporting via user surveys, collected data on user's interpretability perceptions. Statistical analysis for the proposed hypotheses was performed using correlation and regression techniques with data collected through both AWS billing and usage reports and a structured self-report survey to determine the construct validity of the survey measures. Internal validity was established through controlled experimental/resource configurations and maintenance of consistent experimental conditions.

Practical Implementation and Validation:

To operationalize the proposed Cost Interpretability Framework, controlled experiments were conducted within an AWS Free Tier environment. The purpose of this implementation was not performance evaluation but to observe how beginner cloud activities are translated into billing representations and how these representations support conceptual validation.

● **Controlled Free Tier experiments**

All experimental activities were performed using an AWS Free Tier account through the AWS Management Console. The environment was set up on purpose to be like what a beginner would use when they are learning. It used the default settings and the free resources that are available. There were no expensive jobs run on it. This way we could see how the billing would work when someone is just starting out and trying things.



The tests were, like what a beginner would do when they are first using the cloud. This included starting up machines putting data in storage making databases and using serverless functions. We picked these things because they are what people usually do when they are first starting to use the cloud.

- **Experimental Procedure**

Each service scenario followed a consistent execution procedure:

1. A Free Tier–eligible resource was configured using default parameters.
2. A basic user action was performed (e.g., instance execution, file upload, database initialization, or function testing).
3. The AWS billing dashboard and Free Tier usage interface were monitored following the action.
4. Billing visibility was observed across multiple time intervals to capture update delays and metric representation behavior.

This structured procedure enabled comparable observations across services while maintaining experimental consistency.

- **Billing observation results**

Billing-related data were collected through continuous monitoring of AWS billing interfaces, including the Billing Dashboard and Free Tier usage pages. Observations focused on how system-recorded consumption appeared within user-facing dashboards rather than measuring financial cost values.


The study adopted an observation-based qualitative recording approach in which dashboard states were manually documented after each experimental action. Particular attention was given to:

- differences between performed actions and displayed billing metrics,
- delayed dashboard updates,
- technical terminology used in cost representation, and
- abstraction levels of service-specific billing units.

These observations formed the empirical basis for validating interpretability relationships proposed in the conceptual framework.

• Screenshots

AWS Console environment



EC2 instance running

✓ Successfully initiated starting of i-086c1e4ca5ae2abb1

Instances (1/1) Info

Find Instance by attribute or tag (case-sensitive) All states

Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Availability Zone
research_pract	i-086c1e4ca5ae2abb1	Running	t3.micro	Initializing	View alarms +	us-east-1a

researchpract Info

Objects Metadata Properties Permissions Metrics Management Access Points


Metrics Info

CPUUtilization, EBSReadBytes, EBSReadOps, EBSWriteBytes, EBSWriteOps, MetadataNoToken


Label	Min	Max	Sum	Average	02/28 12:15	02/28 12:10	02/28 12:05	02/28 12:00	02/28 11:55
MetadataNoToken	0	0	0	0	0	0	0	0	0
EBSWriteOps	18	130	301	60.2	18	63.4	39.8	50.2	130
NetworkOut	270B	614B	2.13kB	427B	270B	427B	394B	430B	614B
NetworkIn	270B	606B	1.9kB	379B	270B	312B	336B	372B	606B
EBSReadOps	0	1.4k	1.4k	279	0	0	0	0	1.4k
EBSReadBytes	0	60.4MB	60.4MB	12.1MB	0	0	0	0	60.4MB
EBSWriteBytes	102kB	3.72MB	6.53MB	1.31MB	102kB	1.14MB	678kB	892kB	3.72MB
CPUUtilization	0.14%	0.24%	0.9%	0.18%	0.24%	0.16%	0.19%	0.17%	0.14%




S3 bucket data




RDS instance



Lambda execution



Billing dashboard



• Data tables

Table 1. Service Activity Observations

Services	Action Performed	Observation Interval	Dashboard status
EC2	Instance execution	Runtime monitoring	Low utilization but continuous system-level metrics observed
S3	Data upload (168.1 MB)	Observed over dashboard monitoring interval after upload	Storage metrics reflected ~176 MB bucket size with 529 objects; updates appeared aggregated rather than immediate
RDS	Database creation and allocation	Observed during active instance period	Database instance remained allocated with low CPU (~3.6%) and zero connections, indicating resource allocation despite minimal usage
Lambda	Function creation and test execution	Metrics monitored during execution and shortly after instance activity	Low CPU utilization (~0.15%), minimal network traffic, and background EBS read/write activity observed despite negligible workload

Table 2 Operational Interpretation & Cost Signals

Servi ces	User Action	Dashboard Representation	Interability Issue
EC2	Instance executed for short duration (~41 minutes)	Continuous system metrics and instance-hour based monitoring despite minimal workload	Time abstraction between perceived usage and billing measurement
S3	Data upload (168.1 MB storage)	Aggregated storage size and object count updated periodically	Visibility gap due to delayed metric reflection
Lamb da	Single function test execution	Invocation metrics represented through execution duration and abstract compute units	Metric complexity caused by serverless billing abstraction
RDS	Database instance created but largely idle	Active resource allocation shown despite zero database connections	Resource misunderstanding between allocation and actual usage

Results:

The experimental results confirmed a significant gap between raw resource usage and dashboard representation.

By applying the collected data to the Overall Interpretability Index (OII), the study quantified the beginner experience as follows:

1. Quantitative Analysis of Interpretability

Using the framework formula $OII = (TS + PS + TI) - CLI$, the session yielded a low interpretability score.

- i. **Predictability Score (PS):** While the expected cost for Free Tier usage was \$0.00, the actual cost recorded was \$0.41. This result indicates a high deviation, lowering the PS and supporting H3 (alignment between expected and actual costs improves clarity).
- ii. **Traceability Score (TS):** Charges were categorized under "EC2-Other" and "Virtual Private Cloud" rather than being directly mapped to specific resource IDs or the "research_pract" instance. This lack of direct mapping resulted in a low TS, supporting H2.
- iii. **Transparency Index (TI):** In S3 experiments, 168.1 MB was uploaded, but the dashboard aggregated this into **529 objects** with updates that were not immediate. The delay in "real-time" billing updates hindered the user's ability to see the immediate impact of their actions, supporting H1.
- iv. **Cognitive Load Index (CLI):** The presence of technical billing units such as vCPU-hours, EBS read/write activity, and I/O requests for simple tasks increased the CLI. This high cognitive load confirms H4, as it directly reduces the overall interpretability index (OII).

2. Hypothesis Testing Outcomes

- i. **H1 (Dashboard Transparency): Supported.** Monitoring intervals confirmed that update latencies create a disconnect between action and representation.
- ii. **H6 (Service Specific Ambiguity): Supported.** Services like CloudWatch and EBS showed background activity (e.g., 1.01 KB/s or 11.6 KB/s balances) even when the user-initiated workload was zero. This complexity causes higher ambiguity than core compute tasks.

Discussion:

The quantified OII score highlights that the "Mental Cost Model" in the User Interpretation Layer fails because the Presentation Layer (dashboards) prioritizes enterprise-level metrics over learner-centric clarity. The \$0.41 cost deviation is not just a financial issue but a cognitive one; it increases the perceived financial risk, which as hypothesized in H5 reduces the user's willingness to experiment and learn.

Limitations:

While this study provides a foundational framework for cloud cost interpretability, several limitations must be acknowledged:

1. **Scope of Services:** The empirical study focused primarily on core AWS Free Tier services (EC2, S3, RDS). The billing complexities of advanced services like AWS Lambda (serverless) or SageMaker (AI/ML) were not explored.
2. **Sample Size:** The observations were based on controlled experimental setups rather than a large-scale longitudinal study of diverse user groups.



3. **Regional Variation:** AWS pricing and dashboard latency can vary slightly by region (e.g., us-east-1 vs. ap-south-1). This study was conducted within a single region, which may not account for global variance in billing reflection times.
4. **Temporal Constraints:** The study observed billing cycles over a specific period. Long-term "hidden costs" that accumulate over several months (like snapshot storage or elastic IP idling) may not be fully captured in a short-term study.

Future Work:

To build upon the Overall Interpretability Index (OII), future research should address the following:

- **Cross-Platform Analysis:** Applying the framework to Microsoft Azure and Google Cloud Platform (GCP) to compare how different "Presentation Layers" affect beginner cognitive load.
- **AI-Driven Interpretability:** Investigating how LLM-based "Billing Assistants" could translate technical units (vCPU-hours) into natural language to improve the Traceability Score (TS).
- **Dynamic Framework Validation:** Testing the framework against "Pro-Tier" or paid environments to see if the interpretability gap increases linearly with architectural complexity.
- **Educational Integration:** Developing a "Learner-Mode" dashboard prototype based on the findings of this paper and testing its impact on student experimentation rates in university settings.

Conclusion:

This research addressed the critical "fear of billing" that hinders cloud adoption among beginners. By proposing and testing the Five-Stage Cost Interpretability Framework, we moved beyond anecdotal evidence to a quantitative understanding of the cloud learning experience.

Our findings reveal that the primary barrier to cloud experimentation is not the cost itself—as evidenced by the minimal \$0.41 charge—but the Interpretability Gap between user action and dashboard representation. The high Cognitive Load Index (CLI) and low Predictability Score (PS) confirm that current billing interfaces are optimized for enterprise accounting rather than pedagogical clarity.

The practical importance of this study lies in its roadmap for cloud providers and educators: to foster a "fearless" learning environment, cloud transparency must transition from technical logging to human-centric interpretation. By closing the "gap-to-gain," we can ensure that the next generation of developers is limited only by their creativity, not by the ambiguity of their invoices.

References:

1. Amazon Web Services, Inc. (2025). *From cloud cost to cloud value: Creating a culture of cost transparency and accountability*. AWS Whitepapers. <https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-transparency-accountability/cost-optimization-transparency-accountability.pdf>
2. Deochake, S. (2024). *ABACUS: A FinOps service for cloud cost optimization*. arXiv preprint arXiv:2501.14753. <https://arxiv.org/abs/2501.14753>
3. Feitosa, D., Penca, M.-T., Berardi, M., Boza, R.-D., & Andrikopoulos, V. (2024). *Mining for cost awareness in the infrastructure as code artifacts of cloud-based applications: An exploratory study*. *Journal of Systems*



and Software, 215, 112112. <https://doi.org/10.1016/j.jss.2024.112112>

4. Hamza, M., Akbar, M. A., & Capilla, R. (2023). Understanding cost dynamics of serverless computing: An empirical study. *arXiv preprint arXiv:2311.13242*. <https://arxiv.org/abs/2311.13242>
5. Leitner, P., & Scheuner, J. (2020). Estimating cloud service costs of cloud applications: An experiment. *IEEE Transactions on Cloud Computing*, 10(2), 1266-1279.

Cite This Article:

Ms. Kamble R.S., Ms. Dalvi T.F.F. , Mr. Ponnamm S.B. & Ms. Maurya V. (2026). *Cost Transparency and Billing Interpretability Challenges In Aws Environments: An Empirical Study of Beginner Cloud Usage.* **In** *Educreator Research Journal: Vol. XIII (Issue I)*, pp. 86–99.