



# Hidden Cost of Cloud Abstraction

Zaw Zaw Htwe<sup>1</sup>, Arun Udayasuriyan<sup>2</sup>

<sup>1</sup>Student, Department of BSc IT, Faculty of IT & CS, Parul University, Gujarat, India.

<sup>2</sup>Assistant Professor, Department of MCA, Faculty of IT & CS, Parul University, Gujarat, India.

## How to cite this paper:

Zaw Zaw Htwe<sup>1</sup>, Arun Udayasuriyan<sup>2</sup>, "Hidden Cost of Cloud Abstraction", IJIRE-V7I1-101-104.



Copyright © 2026  
by author(s) and  
Fifth Dimension  
Research

Publication. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>

**Abstract:** Cloud computing provides on-demand computing power over the internet, removing the need for physical infrastructure management. Cloud abstraction allows users to manipulate virtual machines as objects, enabling the process of integrating, launching, and running applications in a way that results in a smooth, uniform, and continuous manner. Cloud abstraction conceals technical details and offers smart options, which improve compatibility, adaptability, and scalability. As a result, cloud platforms have become an attractive option for both research and industrial sectors. With the aid of automated resource optimization layers with effective setups, complex and difficult systems can now be built and integrated with minimal effort. However, this ease of use comes with its own hidden drawbacks. By packaging system details away and relying on non-transparent automated services, we can lose sight of critical connections, and we sacrifice the oversight needed to ensure efficiency and reliability. Abstraction layers can make infrastructure fragile, leading to difficult-to-trace errors or failures and dangerous dependence on single providers. This research explores how ease of use compromises system resilience, proposing principles for more sustainable and user-friendly robust durability.

**Keywords:** Cloud abstraction, system fragility, observability, MTTR, cloud resilience.

## I. INTRODUCTION

The Hidden Cost of Cloud Abstraction: Why Convenience Creates Fragile Systems starts by exploring a central conflict in modern infrastructure: while layers of cloud abstraction are designed to deliver simplicity and speed, they mostly hide vital complexities that leave systems more vulnerable during failures [2][3]. Global spending on public cloud services reached \$670 billion in 2023, highlighting a 19.9% increase compared to the previous year driven by the promise of lower operational costs and quick development [7]. However, this recognized efficiency generally ignores the systemic risks introduced by massively abstracted environments [2]. Large-scale incidents, such as the October 2025 AWS US-EAST-1 DNS outage that cascaded through 142 services, including EC2 and Lambda, disrupting major platforms like Snapchat and Reddit [7]. This incident illustrates how these hidden layers force engineers to investigate during a crisis, searching for answers rather than acting on clear system signals. Abstraction can also lead to harmful “antipatterns,” including bloated tooling ecosystems (such as redundant IaC and observability tools) [1]. Likewise, lift-and-shift migrations that avoid true cloud-native redesigns often degrade system stability by favoring short-term simplicity over long-term insight and adaptability [1]. The research evaluates these trade-offs by linking abstraction to Mean Time To Recovery (MTTR) in outage data, revealing how convenience builds technical debt that remains absent from financial records—reported in industry, showing that nearly 32% of cloud expenditure is wasted due to unmanaged complexity [9], [10]. While abstraction is expected to support a projected \$723 billion in end-user cloud spending by 2025 [11], it can also produce fragile systems unless intentionally managed with transparent architecture and proper monitoring.

Cloud abstraction provides simple interface usages and system flexibility [8]. Abstraction-based design adds a layer of operational opacity, which makes it difficult to detect architectural faults and effective system recovery [2]. Many previous research efforts have focused on system performance, scalability, and problems with cost-related cloud operational environments [1], [8]. In these explorations, cloud abstraction is considered only as an architectural benefit for the abstract level [8]. However, only a small number of studies have highly regarded abstraction as a possible cause of fragility [2]. As a result, limited awareness into the underlying infrastructures and dependence on automation to orchestrate system components become a potential for a longer Mean Time To Recovery (MTTR) than expected and hidden technical fragility created via automation [1], [5]. This has led to an understanding gap in how design decisions made using abstraction influence system reliability and robustness across various cloud platforms. This issue needs further examination. The design layers that use abstraction are developing, i.e., abstraction-based designs are evolving due to the introduction of serverless architectures, container orchestration tools, and multi-cloud integration, which are reshaping cloud system architectures [1], [10]. Over time, this adoption widened both the strengths and hidden challenges of layer-based design approaches.

The aim of this research is to examine cloud abstraction contributions to system failure. It assesses the impact of

abstraction-driven decisions on clarity and persistence in cloud environments. This study aims to measure the trade-offs between mean time to recovery (MTTR) and system transparency. Eventually, this article proposes a sustainable abstraction design pattern.

### II.LITERATURE REVIEW

Cloud platforms use abstraction to simplify infrastructure management. It makes managing global infrastructure simple and enables automated provisioning. Prior research has shown that abstraction can reduce operational load during normal operation. It can limit failure visibility and complicate fault identification. Research on large-scale distributed systems explores how invisible links between different parts of the system delay root cause detection because of tightly connected components that make troubleshooting difficult. These findings encourage a better way to examine the impact of abstraction layers on system observability, establishing the basis of the first research question.

Many researchers have spent years analyzing post-mortem reports and measuring the time taken to fix problems (MTTR). Previous studies explore three main causes for long outages: delayed warnings, confusing diagnostics, and complex web dependencies. The problem is that these studies look at the cloud as a single, simple unit. This approach observes increasing abstraction depth in complex problem diagnosis. The investigation on the relationship between abstraction layers and MTTR on the second question suggests that the more obscure the inner workflow, the longer the Mean Time To Recovery (MTTR)

The system health mainly depends on three core Type of data logs, and traces to diagnose problems across complex networks. Unfortunately, highly simplified cloud services often lock away raw data, making observability tools less effective. Although existing approaches enhance spotting failures, they still lack abstraction-induced opacity. Design transparency is still needed for design strategies. This motivates the final goal of this research: proposing awareness-based guidelines for transparent abstraction.

### III.RESEARCH QUESTIONS

- RQ1. How do cloud abstraction layers influence system visibility and failure diagnosis during cloud service outages?
- RQ2. What relationship exists between the level of cloud abstraction and Mean Time To Recovery (MTTR) in large-scale cloud failures?
- RQ3. Can transparency-oriented abstraction design patterns reduce system fragility and improve recovery outcomes in cloud environments?

### IV.METHODOLOGY

#### Failure Visibility and Diagnostic layers

By referencing publicly accessible cloud outage reports and technical postmortems, this study examines how abstraction-driven decisions affect system transparency and problem diagnosis in cloud environments. To understand operational deficiency in large-scale distributed systems, postmortem-based qualitative analysis is widely recommended for complex cloud environments [1]. This paper discusses cloud service incidents where simplified interfaces obscure root architectural failures. These simplified views delayed problem identification and resulted in limited diagnostic data. Prior studies show that higher levels of virtualization obscure interconnections, making it time-consuming during system outages [2]. If observability signals such as logs, metrics, and distributed interconnections across abstraction layers are transparent, troubleshooting remains efficient [4]. System visibility plays a key role in problem fixing, determines how transparency at each individual abstraction layer affects troubleshooting scenarios. Deficiencies can be categorized into different types of errors based on visibility and obscurity, sorting them into three tiers:

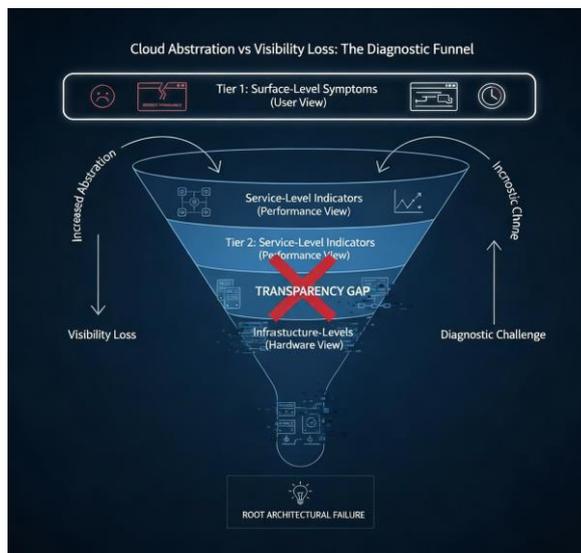


Fig. 1. Conceptual cloud abstraction vs. visibility loss: the diagnostic funnel.

As shown in Fig. 1, abstraction reduces visibility across layers.

- 1) How the user sees it: surface-level symptoms,
- 2) How the service is performing: service-level indicators, and
- 3) What is the underlying hardware doing? : Infrastructure-level signals.

By tracing how each system layer restricts data flow clarity, the study shows that the choices made during the design phase can either help or hinder an operator to figure out and fix problems in the middle of high-stakes cloud outages.

### MTTR Analysis Approach

To analyze the relationship between cloud abstraction and Mean Time To Recovery (MTTR), this study compared postmortem reports of several large-scale cloud service outages to determine whether high level of abstraction slows down the recovery time. Publicly accessible incident reports are analyzed to map out exactly how long it took to spot the problem, diagnosis duration, time to response (TTR), and operational normalization [1]. By using these timestamps, we can clearly measure where the system's complexity influenced the speed of the recovery. Failures are sorted based on how much of the system was hidden from users, ranging from traditional servers to fully automated serverless platforms. MTTR statistics are evaluated across these categories to find patterns to see if diagnostic visibility or hidden links are the cause. This approach prevents the research from jumping to conclusions and makes it accessible how architectural design decisions and repair speed are actually linked.

### Transparency-Oriented Design Pattern Analysis

To investigate whether system transparency can reduce system fragility and improve recovery timelines, this study analyzes technical architecture and past cloud system failure reports by adopting a qualitative

Design-oriented analysis. The analysis focuses on identifying designs that expose moderate internal state, dependencies, or failure signals to engineers without sacrificing simplicity and abstraction [2], [4]. Publicly available system design documentation, engineering blogs, and incident postmortems are extracted to propose a set of design patterns that focus on opacity rather than secrecy, such as creating multi-level data views that work for both high-level and deep-level issues, clearly mapping out interrelated connections, and building a system with self-failure reporting capabilities. These design patterns are compared to real-world system crashes in which hidden layers limited opacity and contributed to system failures [2]. These comparisons do not just focus on recovery speed but on how design helps an engineer make the right decisions during outages. The design patterns are evaluated on how they impact system comprehensibility, fault detection, and recovery timelines. By contrasting transparent and encapsulated methods, the study examines how these choices affect system resilience and fragility, providing Design-level insights to make smarter design choices for long-term rather than prescriptive implementation guidelines. To answer the question, transparent abstraction design patterns can reduce system fragility by making the system much more resilient because it reveals complex architecture that usually hinders problem identification. It can significantly lower Mean Time To Recovery (MTTR). This ensures abstraction patterns re-expose system failure signals, improving human and automated response.

## V. ANALYSIS AND DISCUSSION

This section breaks down insights from cloud outage postmortems, recovery timelines, and design-oriented evaluation. The discussion is structured according to the three research questions, focusing on visibility, service restoration time, and the design of cloud abstractions.

### Abstraction and Failure Visibility (RQ1)

Based on the cloud service incident reports, higher levels of abstraction often reduce observability during system failures [2]. Engineers were only exposed to high-level symptoms, such as service connectivity loss or increased error rates, whereas the root causes remained hidden behind a simplified dashboard. This made effective diagnosis delayed, since teams had to inquire about the system's internal behavior without having infrastructure-level signals. Not having enough transparency in the system design caused delays in fault identification and increased manual intervention during failures. These points show that abstraction layers serve as a primary factor to influence both the speed and precision of error detection.

### Abstraction and Mean Time to Recovery (RQ2)

Comparing various outage timelines shows that incidents with highly abstracted environments generally tend to have longer Mean Time To Recovery (MTTR). Abstraction simplifies system operation but limits clarity when troubleshooting cloud failures, which impacts recovery time [1]. Systems that provide transparent architecture across both service and underlying infrastructure enable faster fault detection and efficient recovery.

### Transparency-Oriented Abstraction design Patterns (RQ3)

The design-oriented evaluations indicate that abstraction itself is not the thing that makes system resilience fragile. The real problem arises when abstraction masks critical diagnostic signals [2]. Focusing on transparent design patterns helps by selectively exposing internal state, component dependencies, and error data. To answer the question, transparent abstraction design patterns can reduce system fragility by making the system much more resilient because it reveals complex architecture that usually hinders problem identification. It can significantly lower Mean Time To Recovery (MTTR).

## VI. THREATS AND VALIDITY

This study is based on public postmortems and incident logs, which means the depth of analysis is limited by what is published in those documents. Publicly available documents are usually written after the problem has been resolved and may not include internal decisions and small, mid-recovery steps during service restoration. Consequently, the timestamps during system outages might be approximations based on the available narrative rather than exact.

There is also another challenge in how to label abstraction levels. Cloud platforms often combine multiple layers and categorize them into distinct groups, requiring interpretation. Even though there are reasonable sorting processes, some classification bias may still exist. Additionally, the analysis looks specifically at large-scale public cloud outages; the findings might not perfectly generalize to private or on-premise environments.

This study focuses on analyzing actual timelines and qualitative observations instead of using a controlled experiment. Despite these limitations, the incidents provide enough evidence to achieve the goals of this initial investigation.

## VII. CONCLUSION

This study explored the impact of cloud abstraction on system fragility during failures, focusing on three main areas: visibility, diagnosis, and restoration timeline. The findings suggest that while cloud abstraction makes systems easier to manage, it can also hide internal behavior. On the contrary, by masking internal architecture, abstraction adds a layer to the resolution cycle, making fault identification and recovery more difficult. By reintroducing transparent abstraction design, a cloud system may retain simplicities and improve system resilience. Overall, this work highlights the importance of simplicity and visibility, providing insights into how cloud platforms can be structured to support a quicker and more reliable restoration cycle during failures.

## Reference

1. B. Beyer, C. Jones, J. Petoff, and N. R. Murphy, *Site Reliability Engineering: How Google Runs Production Systems*. Sebastopol, CA, USA: O'Reilly Media, 2016.
2. R. I. Cook, "How Complex Systems Fail," *Cognitive Technologies Laboratory*, Univ. of Chicago, Tech. Rep., 2000.
3. J. Allspaw, "The Infinite Hows," *USENIX; login*: vol. 43, no. 4, pp. 26–32, 2018.
4. C. Sridharan, *Distributed Systems Observability*. Sebastopol, CA, USA: O'Reilly Media, 2018.
5. D. Yuan, Y. Luo, X. Zhuang, G. Ren, Z. Li, X. Zhang, B. Jain, and M. Stumm, "Simple Testing Can Prevent Most Critical Failures: An Analysis of Production Failures in Distributed Data-Intensive Systems," in *Proc. 11th USENIX Conf. Operating Systems Design and Implementation (OSDI)*, 2014, pp. 249–265.
6. Google Cloud, "Google Cloud Incident Reports," [Online]. Available: <https://status.cloud.google.com/incident-history>. Accessed: Jan. 2026.
7. Amazon Web Services, "AWS Post-Incident Summaries," [Online]. Available: <https://aws.amazon.com/premiumsupport/technology/pes/>. Accessed: Jan. 2026.
8. D. L. Parnas, "On the Criteria to Be Used in Decomposing Systems into Modules," *Commun. ACM*, vol. 15, no. 12, pp. 1053–1058, Dec. 1972.
9. Flexera, "2023 State of the Cloud Report," Flexera, Itasca, IL, USA, 2023. [Online]. Available: <https://www.flexera.com/blog/cloud/cloud-computing-trends-2023-state-of-the-cloud-report/>
10. Google Cloud, "2023 Accelerate State of DevOps Report," DORA Research, 2023. [Online]. Available: <https://cloud.google.com/devops/state-of-devops>
11. Gartner, "Gartner Forecasts Worldwide Public Cloud End-User Spending to Reach \$723 Billion in 2025," Gartner Press Release, Oct. 2024. [Online]. Available: <https://www.gartner.com/en/newsroom>