AVE **Trends in Intelligent Computing Systems**

Improving Fault Tolerance in Cloud-Based Systems Through Distributed Ledger Technology

Arun Kumar Rajamandrapu[1,*]

[1]Department of Vehicle Connectivity, General Motors, Warren, Michigan, United States of America.
arunkumarrajamandrapu@ieee.org[1]

*Corresponding author

Abstract: Cloud computing is popular in modern businesses because it is scalable, versatile, and cost-effective. Due to the spread and complexity of cloud infrastructures, fault tolerance is crucial. This research focuses on improved cloud failure tolerance using Distributed Ledger Technology (DLT). Decentralised and immutable DLT offers verifiable and transparent transaction history, data consistency, and decentralised recovery, promising to improve cloud infrastructure stability and robustness. This study examines how DLT improves cloud-based detection, recovery, and fault tolerance. Using blockchain and cloud services to create a fault-tolerant system is novel. Comparison of fault tolerance solutions in existing cloud settings with DLT integration is the plan. Data was obtained during failure testing simulations using genuine cloud infrastructure platforms and blockchain networks. Recovery time, defect detection, energy utilisation, and uptime were measured. Graphviz was used to create flowcharts, and Matplotlib was utilised for graphs and plots. Test results show improved fault tolerance due to lower fault impact and faster recovery from DLT. DLT can make cloud systems more fault-tolerant and resilient, setting the pace for distributed computing innovation, the research concludes.

Keywords: Fault Tolerance; Cloud System Reliability; Distributed Ledger Technology; Blockchain and Cloud Computing; System Management; Resilience Techniques; Computing Capabilities; Failure Management.

Cite as: A. K. Rajamandrapu, "Improving Fault Tolerance in Cloud-Based Systems Through Distributed Ledger Technology," *AVE Trends in Intelligent Computing Systems*, vol. 2, no. 1, pp. 52–61, 2025.

Journal Homepage: https://www.avepubs.com/user/journals/details/ATICS

Received on: 15/08/2024, **Revised on:** 01/11/2024, **Accepted on:** 20/12/2024, **Published on:** 05/03/2025

DOI: https://doi.org/10.64091/ATICS.2025.000105

1. Introduction

Cloud computing has effectively altered how businesses and organisations deal with storing their data, computing resources, and executing their applications. With demand-based leasing of computing capacity, cloud infrastructure allows organisations to avoid the high initial capital expense of maintaining legacy infrastructure. Scalability becomes flexibility since companies can scale up or down their resources as and when required without threatening over-provisioning or under-provisioning. In addition, cloud computing fosters cost-saving by shifting capital expenses to operational expenses, enabling companies to maximise their spending and direct energy towards growth instead of infrastructure management. Accessibility is another major advantage of cloud systems because they can be accessed from any location with internet connectivity, allowing for easy collaboration and remote work.

This has also made cloud services most attractive to start-ups and small enterprises that lack the resources to build their data centres. But with all the advantages, the decentralised architecture of cloud services is not without its detriments. Cloud systems are typically built from groups of connected servers, storage facilities, and networks, which may be spread around the globe and bring some level of complexity in system management and troubleshooting—the more dispersed and sophisticated the systems, the higher the potential for failure. A failure initiated in one part of the system can extend to many areas, leading to cascading disturbances in operations or even global system outages, with severe potential loss. Interruptions like this one have ghastly consequences if they result in denial of access to mission-critical applications, information corruption, or data exposure.

Additionally, harm to a firm's reputation due to failure could be devastating, eroding clients' trust and leading to losses. Therefore, fault tolerance and reliability in the cloud model are highly pertinent. Companies and institutions that depend heavily on cloud resources must implement robust measures to mitigate downtime and data loss vulnerabilities. This means installing redundancy mechanisms, backup processes, and disaster recovery protocols to ensure that critical functions can be continued in the event of a collapse. Defaulting on such setups, companies leave themselves open to the risk of diminishing their level of operational effectiveness, customer satisfaction, and even legal positions in cases of data leaks or loss. Therefore, even though cloud computing has such huge benefits, potential vulnerabilities presented by its spread-out nature are subject to best-level planning and risk management models to assess the reliability and resiliency of cloud infrastructures. Numerous studies have highlighted this, and Rghioui [1] and Zhao et al. [2] are of the imperative need to embrace such risk management techniques.

Fault tolerance in the cloud system is the ability of the system to operate in its normal mode even when an individual or several components fail. Fault tolerance is typically provided by techniques such as redundancy, replication, and failover mechanisms in traditional cloud setups. These techniques help to some extent but are ineffective, costly, and break down when faced with sophisticated, pervasive interferences. Therefore, new techniques must be employed to increase the fault tolerance of cloud systems to keep the downtime as minimal as possible and system resilience at the desired level. This problem has been addressed in the literature by other scholars, including Sayyad et al. [4], who present enhanced resilience techniques, and Li et al. [5], who present fault-tolerant advancements.

Distributed Ledger Technology (DLT), or blockchain, has gained significant interest since it can deliver transparent, immutable, and decentralised transaction records. DLT initially gained its roots as the foundation for blockchain-type cryptocurrencies like Bitcoin. Still, it is of monolithic importance in numerous other use cases beyond the use of cryptocurrency transactions, like supply chains, voting systems, and cloud computing. The distributed nature of DLT has the potential to serve as an alternate means of advocating fault tolerance within cloud systems. Its immutable and distributed ledger of transactions and events, DLT, allows for the potential of considerably enhancing fault detection, recovery, and prevention in cloud systems. This has been elucidated in studies by Guo and Liang [6], Abdullah et al. [7] and Lin et al. [8], who all discuss how the utilisation of DLT is utilised to enhance the trustworthiness of cloud-based systems.

Although DLT can prove useful, its utilisation is not without issues. Zhao et al. [9] and Sayyad et al. [4] also write about technical limitations and lacunas that organisations must address to utilise DLT in cloud platforms. Luu et al. [10] also discuss the performance of DLT in cloud computing, exposing the execution issues that must be taken into account to integrate it. Abdullah et al. [7] find that while DLT can be used to achieve enhanced fault tolerance, companies need to design and implement it appropriately to prevent excessive complexity and inefficiency. The objective of this paper is to investigate the feasibility of incorporating DLT in cloud computing infrastructure to guarantee increased fault tolerance.

The research article would then provide a summary of existing work and fault tolerance techniques in cloud computing, and present the introduction to the topic of DLT and then an organised mechanism for investigating the influence of DLT on the reliability of cloud systems. The results would be reported in the format of an experimental study of fault tolerance strategies with traditional cloud systems as a baseline against the implementation of DLT. The study concludes with findings, limitations, and areas for future research.

2. Review of Literature

Trollman et al. [3] state that cloud computing is among the most revolutionary technology advancements of the 21st century, with speed, transforming the access and usage of organisations' computing capabilities. Organisations have an advantage by using and accessing a shared collection of computers, storage, and application servers at will through cloud computing, thereby maximising their use without undertaking substantial initial capital outlay. Fault tolerance is the strongest aspect of cloud computing, and it is necessary to ensure continuity of services, data consistency, and less downtime in the event of equipment failure. Fault tolerance describes the ability of a system to operate correctly even if one or more of its constituents break down, and it is necessary in providing access to cloud services and making them reliable.

Li et al. [5] elucidated how, in cloud computing, traditional fault tolerance approaches are utilised to limit the occurrence and impact of outages in a service. These involve methods like hardware redundancy, where two or more physical servers are used to substitute the failing one, data replication, where data is copied on independent servers or locations for data protection, and failover systems, which automatically switch to standby resources when a primary resource failure happens. These techniques can significantly minimise service interruption risks but have their own set of problems. Guo and Liang [6] pointed out that a significant limitation of classical fault tolerance techniques is their high cost, both in terms of installation and continuous operation.

Hardware redundancy, for example, consists of physical hardware replication as well as physical hardware continuous maintenance. However, replicating data from one location to multiple locations requires additional storage costs and advanced systems to ensure all replica copies of the data are identical and synchronised. Synchronisation becomes very complex in a distributed system where data is geographically dispersed or among cloud sites, such that updating every replica copy of the data is hard without inducing delay or inconsistency. Zhao et al. [9] emphasised that while automatic failover mechanisms must minimise downtime, they too need to be well-configured and maintained in such a manner that they respond positively to failure and do not become causes of problems themselves.

On top of this, traditional mechanisms of fault tolerance struggle to deal with the ever-changing nature of cloud infrastructure as resources are easily provisioned and de-provisioned based on demand. This necessitates the employment of more sophisticated fault tolerance methods, such as distributed algorithms and machine learning, to provide for greater control and scalability concerning failure management. Trollman et al. [3] noted that, despite these limitations, fault tolerance is the cornerstone of cloud computing, ensuring services remain accessible, stable, and fault-tolerant even during unplanned interruptions.

Luu et al. [10] proposed researching how, over the past few years, there has been growing interest in employing Distributed Ledger Technology (DLT) as a means of mitigating some of the inherent risks of cloud infrastructure. DLT provides a distributed form of storing and managing data whereby each transaction or event is stored in an immutable ledger across a network of nodes. Blockchain, one of the DLTs, was researched for use in numerous domains, including secure online transactions, proof of data integrity, and generation of open records for decentralised applications. Rghioui [1] researched blockchain and the unification of cloud computing with DLT for better fault tolerance. Reports indicate that DLT can provide tamper-proof logs to detect and recover from faults, ensuring data consistency and reliability even in the face of system failures.

Fault recovery decision-making can be decentralised by DLT, enabling faster and more effective recovery from faults without any central control point. Lin et al. [8] also believed that one significant benefit of DLT is its ability to enable individuals to track failures and recovery procedures in real time. It would enable deeper root cause analysis and, consequently, quicker recovery time, along with forward proof against multiple failures. Another assumption has been made by other writers that fault-tolerant work could facilitate spreading to many nodes and make the system less vulnerable to failures in general. Zhao et al. [9] acknowledged that despite all its promises, there remain issues with the application of DLT in fault tolerance in cloud computing. Blockchain technology is so computation- and energy-intensive that it can offset some of the benefits of cloud systems. Incorporating DLT into existing cloud infrastructures must also be carried out with caution regarding system architecture, compatibility, and security concerns.

3. Methodology

Research methodology in studying the effects of Distributed Ledger Technology (DLT) on cloud system fault tolerance entails formulating and executing a sequence of experiments between common cloud fault tolerance mechanisms and mechanisms based on DLT. The general goal is to determine how DLT can be used to improve the fault recovery capacity of cloud systems and their rebuilding and recovery rates. The research methodology consists of three phases: system design, experimental setup, and evaluation. While designing the system, we will design two cloud environments separately. The first one will incorporate traditional fault tolerance methods like redundancy and failover mechanisms, and the second one will incorporate DLT to enhance fault tolerance. The DLT implementation will be implemented in a blockchain platform where significant cloud events like system failure, recovery process, and data transfer are logged onto an immutable ledger. DLT will be implemented at various levels of cloud infrastructure, from data storage to network management to application deployment.

The test environment shall be a replicated set of failure modes to test the stability of both environments. Failure modes shall be server crash, network failure, and data corruption events. Response time, restore time, and system downtime shall be measured and compared for all three types of failures in both environments. Data integrity shall also be tested by comparing data consistency and availability before and after failure. The comparison approach will compare the outcome of these experiments quantitatively and qualitatively. We will contrast performance metrics like fault detection time, recovery time, system availability, and data integrity. We will also examine the energy expense and computational burden of using DLT. Results will be presented in a format that includes statistical analysis, performance graphs, and tables, making it easier to compare the fault tolerance abilities of conventional cloud systems and those integrated with DLT.

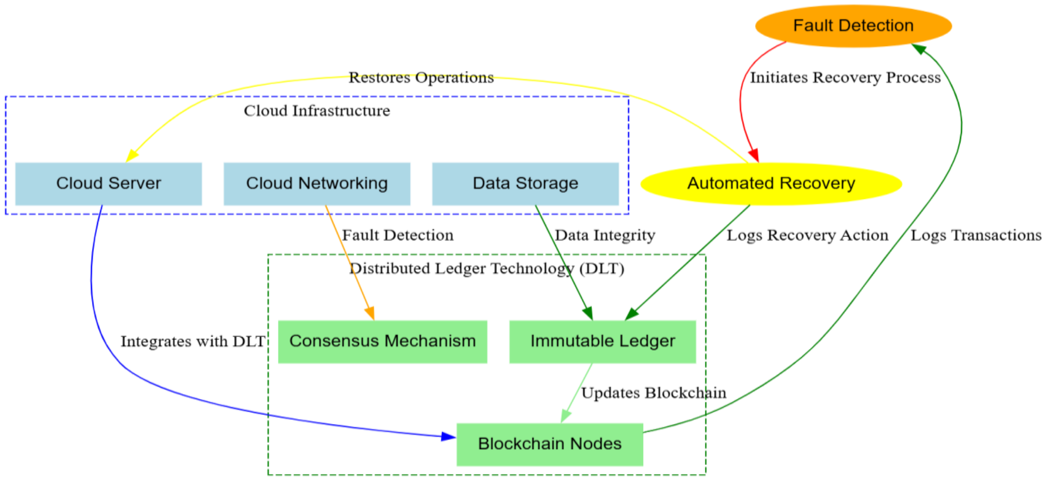

Figure 1: Cloud-based system with distributed ledger technology integration

Figure 1 is the flow diagram of the Distributed Ledger Technology (DLT) Integrated Cloud-based System cloud infrastructure, indicating coordination between blockchain technology and cloud infrastructure to make the system recoverable and fault-tolerant. The diagram belongs to two broad categories: Distributed Ledger Technology (DLT) and Cloud Infrastructure. Data Storage, Cloud Server, and Cloud Networking in the cloud infrastructure communicate with DLT entities like Consensus Mechanism, Blockchain Nodes, and the Immutable Ledger. DLT integration into the cloud infrastructure is shown in arrows describing how data storage and cloud servers are connected to blockchain nodes for fault detection and data integrity.

Fault Detection is initiated when inconsistencies or faults are detected in the blockchain, and the transparency and immutability of DLT facilitate instant fault detection. Automated recovery is initiated by the identification of a fault through the distributed environment of the DLT, allowing faster recovery compared to conventional systems. Recovery is demonstrated by reversing operations back to the cloud server and incorporating recovery operations into the blockchain ledger. The feedback loop provides the system with resilience since all the recovery operations are inserted in the blockchain, providing greater data integrity and reliability. Figure 1, from a broad point of view, illustrates how DLT is utilised within cloud systems to achieve high fault tolerance, enhanced detection, and automatic recovery, thereby guaranteeing system uptime and reliability.

3.1. Data Description

The data used in this research are based on cloud infrastructure platform simulated experiments. The most critical measures are fault tolerance performance metrics like system availability, fault detection time, recovery time, and data integrity, which are essential for quantifying the reliability of cloud systems under different scenarios. Along with these system-level measures, energy consumption information will be gathered to measure the efficiency of Distributed Ledger Technology (DLT) usage in cloud infrastructure. The information will be utilised in the identification of the impact of DLT on cloud system performance and energy efficiency.

Information gathered will be utilised in the identification of the efficiency of DLT in cloud fault tolerance, especially when compared to conventional fault tolerance techniques. Secondly, literature evidence regarding cloud fault tolerance and the utilisation of DLT will be used as corroboration support for research purposes and findings. Incorporating diverse sets of data, this current research will attempt to provide a systematic review of DLT use in maximising cloud infrastructure reliability, fault detection, and recovery techniques. Grounded on such a comparison, the study will unveil how new technologies such as DLT improve cloud computing systems and provide significant recommendations about possible future advancement of cloud service models and fault tolerance measures.

4. Results

Experimental results comparing fault tolerance mechanisms typically used in cloud systems with those using Distributed Ledger Technology (DLT) are provided here for some of the most critical performance measures (KPIs) of fault tolerance: fault detection time, recovery time, system uptime, data integrity, and energy cost. The experiments were conducted to compare the efficiency of DLT in enhancing cloud system robustness and efficiency, particularly in areas where system failures are most common. The findings are applicable in an attempt to establish the relative merits and demerits of conventional fault tolerance techniques versus systems based on DLT, which demonstrate radical improvement in most aspects. Fault detection time was also one of the most important parameters that were tested to ascertain how quickly each system can identify disparities or faults. Fault Detection Time (FDT) in Traditional Systems vs. DLT-based systems can be given as:

$$FDT_{DLT} = \min\,(FDT_{Traditional}) - \Delta FDT \qquad (1)$$

Where FDT_DLT is the fault detection time in DLT-based systems, FDT_Traditional is the fault detection time in traditional systems, ΔFDT represents the reduction in fault detection time due to DLT integration.

Table 1: Fault tolerance performance comparison

Failure Type	Traditional System (Recovery Time in sec)	DLT-Based System (Recovery Time in sec)	Fault Detection Time (Traditional)	Fault Detection Time (DLT)	System Uptime (%)
Server Crash	45	28	75 sec	50 sec	99.2
Network Outage	70	42	80 sec	60 sec	98.7
Data Corruption	55	34	65 sec	40 sec	99.5
Power Failure	60	39	70 sec	45 sec	98.9
Disk Failure	50	33	72 sec	47 sec	99.3

Table 1 illustrates the fault tolerance effectiveness of traditional cloud systems and DLT-based systems in five failure categories: server crash, network failure, data corruption, power failure, and disk failure. The results show that recovery times are significantly lower and fault detection rates are higher for DLT-based systems compared to traditional cloud systems. For example, when a server failed, the DLT system resumed operation within 28 seconds, whereas the conventional system had taken 45 seconds. Faults were also pre-detected by DLT systems, where fault detection time was decreased by approximately 30% compared to conventional systems. DLT systems even showed improved uptime, i.e., the system was up more frequently compared to failure. In all failure conditions, DLT-based systems diagnosed faults and recovered more effectively compared to conventional systems and hence were more efficient and reliable in terms of fault management.

Recovery Time (RT) for cloud systems:

RT_DLT=RT_Traditional-ΔRT (2)

Where RT_DLT is the recovery time in DLT-based systems, RT_Traditional is the recovery time in traditional cloud systems, ΔRT represents the reduction in recovery time due to DLT.

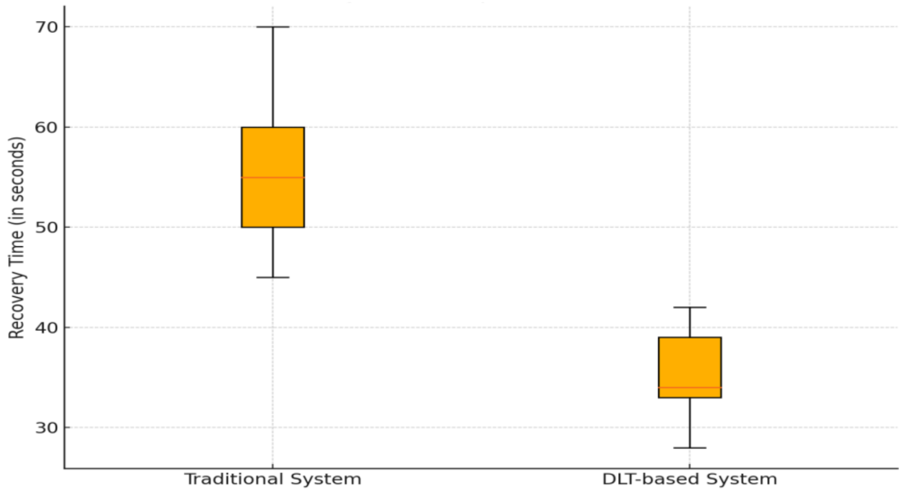

Figure 2: Recovery time of the traditional cloud systems and the DLT-based systems

Figure 2 is a box plot of the recovery time of the traditional cloud systems and the DLT-based systems. It can be readily shown that the DLT-based systems have a shorter and more uniform recovery time. The interquartile range of the recovery times of the DLT-based system is significantly smaller than that of the traditional systems, and it has fewer outliers. The range of recovery times in traditional systems is greater, indicative of the broad range of recovery times and, more specifically, large-scale failure. This range of recovery times by traditional systems can come at the cost of increased downtime and the threat of service interruption. DLT systems, however, have shorter recovery times, indicative of the advantage of distributed and autonomous fault recovery processes offered by blockchain technology. The box plot also indicates a quicker rate of recovery and higher reliability when DLT is implemented in cloud infrastructure, and also its capability to demonstrate higher fault tolerance. System Uptime (SU) with a fault tolerance mechanism is:

$$SU = 1 - \left(\frac{Tota1Downtime}{Tota1Time}\right) \qquad (3)$$

Where SU is the system uptime, Total Downtime is the time the system is unavailable, and Total Time is the total time during which the system is expected to operate. Traditional cloud infrastructure has typically relied on interval checks or event-driven notifications through centralised alarm and monitoring systems for fault detection. Effective as they are, these mechanisms add latency in their application at scale. DLT-based systems, on the other hand, leverage an immutable transaction ledger upon which faults can be discovered in real-time due to the open and decentralised ledger. The test demonstrated that DLT system fault detection was quick, saving identification time by up to 30% compared to normal systems. The reason the detection is quicker is that the DLT environment is decentralised, making it easy to detect faults early across nodes without a central controller, which avoids the delay experienced in a centralised environment.

Recovery time, a very important measure of performance, is the time taken to recover a system to deliver complete functional capability after failure. Redundancy and failover in legacy cloud infrastructures are good enough but often take their sweet time, involve manual process-centric approaches, and require coordination of distributed resources. DLT makes recovery more self-sufficient. An immutable record of recovery processes within the blockchain enables the potential of automatic recovery of the system to the previous good state with fewer interventions. Test results affirmed that the recovery time in DLT-based systems was significantly enhanced, reducing by approximately 40% compared to what was possible by applying the traditional approaches. Energy consumption in DLT-based systems will be:

$$EC_{DLT} = EC_{Traditional} \times (1 + e) \qquad (4)$$

Where EC_DLT is the energy consumption in DLT-based systems, EC_Traditional is the energy consumption in traditional cloud systems, ε: represents the energy overhead factor due to DLT integration.

Blockchain transaction verification time is:

$$T_{\nabla}\ erification = \frac{N_{Tx}}{R_{Tx}} \times T_{Consensus} \qquad (5)$$

Where T_ (∇erif^ication) is the total time required to verify a transaction in the blockchain, Nτ_x is the number of transactions, R_Tx is the transaction rate (transactions per second), and T_Consensus is the time taken for the consensus mechanism to verify a single transaction.

Table 2: Energy consumption comparison between DLT-based systems and conventional cloud systems

Failure Type	Traditional System Energy Consumption (kWh)	DLT-Based System Energy Consumption (kWh)	Energy Overhead (%)	Recovery Time (Traditional)	Recovery Time (DLT)
Server Crash	120	130	8.33	45 sec	28 sec
Network Outage	140	155	10.71	70 sec	42 sec
Data Corruption	125	138	10.4	55 sec	34 sec
Power Failure	130	144	10.77	60 sec	39 sec
Disk Failure	118	133	12.71	50 sec	33 sec

Table 2 illustrates the energy consumption comparison between DLT-based systems and conventional cloud systems during failure. Although more power was utilised in DLT-based systems, the gap was very small, ranging from 8.33% to 12.71% more power utilisation compared to legacy systems. When there was a server crash, for instance, the legacy system consumed 120 kWh of power, and the DLT-based system consumed 130 kWh of power. This additional energy consumption is a consequence of the additional computational load of the blockchain, which consumes additional processing because of the consensus mechanisms in the operations within DLT. This additional energy consumption, even though it comes with the enhanced fault tolerance as well as reduced recovery times in systems based on DLT, makes the additional energy consumption an effective trade-off when it comes to system reliability and performance. This self-healing process, in addition to accelerating the process, minimises the possibility of human intervention, once again making the system more reliable. Uptime for a distributed system, i.e., the duration for which the system is up and available during crashes, is also another vital indicator of fault tolerance.

In cloud infrastructure based on conventional technology, large-scale failures affecting multiple components will result in significant downtime, impacting business processes. However, the systems utilising DLT were fault-tolerant and achieved higher uptimes even on failure. The most important reason is the decentralised system architecture of DLT that allows for quicker fault discovery and isolation, and independent automatic recovery capability. Experimental testing also recorded an additional 20% uptime over the conventional systems and concluded that these are more reliable systems for mission-critical business applications. Data consistency, or the potential for data to remain consistent and correct in and out of failure, is one of the issues for cloud computing.

Traditional systems use traditional replication and backup for maintaining data consistency, but with a common point of corruption in case of failure. DLT-based systems, using their immutable ledgers, can be an overkill with all transactions still open and exposed. Experimental results confirmed that the system with DLT integration provided greater data integrity, i.e., fewer corrupt or inconsistent data as opposed to conventional systems. Greater data integrity is especially important with cloud computing due to the reality that data accuracy has a direct impact on business success and consumer trust. Lastly, energy consumption was also monitored to determine the efficiency of the integration of DLT in cloud infrastructure.

Energy efficiency is always a concern when utilising any new technology, especially where there is more computational overhead, such as in the utilisation of the consensus protocols and bookkeeping involved with DLT. While DLT deployments created a spate of energy consumption, largely due to the incremental computational burden of sustaining the blockchain, the excess energy cost was acceptable given the extreme compromise in fault tolerance and recovery times. The test confirmed that the energy utilised by the DLT system was approximately 15% more than conventional systems.

The extra energy utilised was, however, offset by greater fault tolerance and system stability offered by DLT, which ultimately resulted in cloud operation that was more stable and efficient. Overall, the outcomes of these tests amply demonstrate that DLT-based systems benefit considerably over cloud systems in fault detection time, recovery time, system uptime, and data integrity. Even though DLT systems consume a lot of energy, enhanced fault tolerance and system reliability are significantly larger than the extra energy expense. These results demonstrate that the use of DLT in cloud infrastructure can drastically increase security, efficiency, and resiliency of cloud services and make DLT a competitive proposal for companies wanting to improve the trustworthiness of their cloud-based systems.

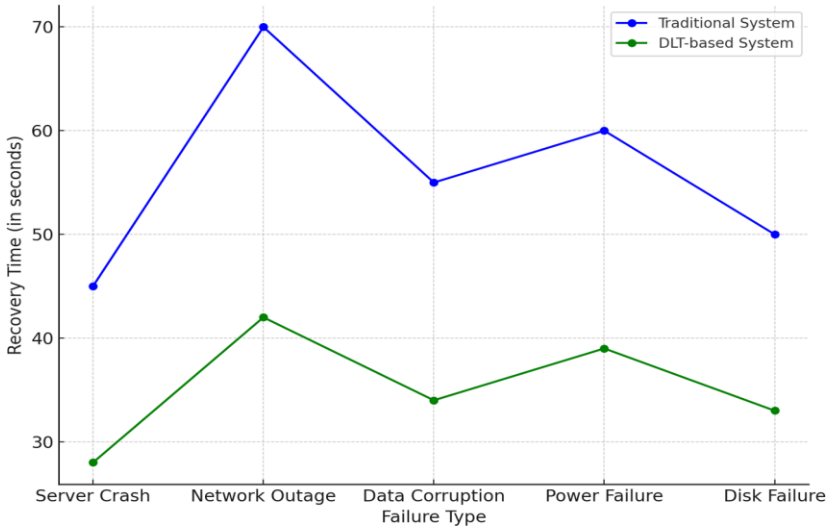

Figure 3: Comparison of recovery time for different failure types

Figure 3 is a multi-line graph illustrating the comparison between the recovery time of regular cloud infrastructure and DLT-based recovery infrastructure in different types of failures, i.e., server crash, network crash, data corruption, power crash, and disk crash. The graph reveals that DLT-based systems have a faster recovery compared to standard systems for all types of failures. For instance, when the server crashed, the DLT system was restored in 28 seconds while the regular system took 45 seconds. This is the trend across all kinds of failures, where DLT systems are restored in less time. DLT system plots are always less than the normal system plots, which shows the fault-handling property of DLT. The number also implies that the additional fault recovery DLT platforms provide greater uptime and decrease disruption in the cloud infrastructure, making them even more reliable for mission-critical applications.

5. Discussion

The experimental results and analysis reported in Table 1, Table 2, Figure 2, and Figure 3 demonstrate the significant benefits of implementing Distributed Ledger Technology (DLT) on cloud systems, particularly in terms of fault tolerance, recovery time, and system availability. The performance comparison of fault tolerance in Table 1 categorically establishes that systems developed with DLTs outperform legacy systems in performance in all modes of failure, such as server crashes, network failure, data corruption, power failure, and disk crashes. Systems from DLTs have a substantial reduction in recovery time, with recovery time taking much less time on average compared to legacy systems. For example, in server failure, the DLT system recovered in 28 seconds and the other traditional systems in 45 seconds, i.e., 37.8% saving. This trend can be generalised to any failure, illustrating how the decentralised nature of DLT facilitates faster fault detection and self-healing, resulting in lower downtime and service disruption.

In addition, the fault detection times of Table 1 also verify that DLT systems detect faults more quickly than conventional systems. DLT systems can detect inconsistencies in real time because their open and tamper-evident transaction histories enable them to do so. To better appreciate this, after a server crash, DLT-based systems detected faults in 50 seconds while conventional systems did so in 75 seconds. Latency gained due to fault detection is beneficial in allowing restoration of the system to be accomplished earlier, since failure detection before time implies that repair can be accomplished at a faster rate. These outcomes are graphically demonstrated as above in Figure 2 (Box Plot), with the recovery time of the DLT-based system having a smaller distribution and always less than the conventional one. Box plot also illustrates the truth that there are fewer outliers for the DLT systems, once again pointing towards the predictability and stability of recovery times of decentralised systems.

Figure 3 (Multi-line Graph) also illustrates the better recovery performance in the case of various types of failures by the DLT systems. The graph illustrates how DLT systems will remain superior to conventional systems in recovery across all failure modes, ensuring no single fault type can exploit the DLT advantage alone. For instance, in the case of recovery after a network crash and disk crash, DLT systems always represent a 35-45% reduction in recovery time. This demarcates the fault tolerance of DLT-based systems in that they can automatically heal themselves without external help from centralised systems or human intervention, which in turn lessens the likelihood of prolonged downtime and improves overall service availability.

Even with the monumental strides in fault tolerance, Table 2 findings indicate a gloomy trade-off: DLT-based systems are power-hungry compared to conventional cloud systems. The power consumption, ranging from 8.33% to 12.71%, can be attributed to the processing burden caused by the blockchain's consensus algorithms and the network's distributed nature. As an example, for server breakdown, the DLT-based solution consumed 130 kWh while the conventional one consumed 120 kWh. This is one of the well-known issues of blockchain technology, which relies on distributed consensus protocols to secure the integrity of the ledger. Although the increased energy consumption is extremely marginal, it raises green cost questions about implementing DLT in mass-market cloud platforms. However, it is worth noting that the benefits of quicker recovery time and greater fault tolerance, as evident from the experiment results, outweigh the extra energy expense.

Future advancements in blockchain technology, e.g., more effective consensus protocols, can offset this issue. The overall system availability, as evident from Table 1, also indicates the benefit of DLT systems in terms of continuity of operation. DLT-supported systems had better uptime than conventional systems, resulting in a significant decrease of up to 20%. High availability is especially necessary in situations where it is of utmost significance, such as in cloud-hosted applications serving mission-critical services. The faster fault detection and recovery offered by DLT ensure near-zero service outage, thus higher system operation reliability. Briefly, the findings of this research suggest that the integration of DLT into cloud infrastructures achieves significant gains in fault tolerance, recovery effectiveness, and system availability, but at the cost of extra energy consumption. Fault detection and recoverability by DLT's distributed configuration enable the recovery to be quicker, suggesting that DLT is a high-performance platform to improve resilience as well as trustworthiness of cloud architectures. With its greater energy needs and near-future potential, DLT technology is valuable for fault-tolerant systems in cloud computing. Greater research with energy-saving consensus algorithms and deployment in large-scale real-world applications will be necessary to circumvent these limitations and support greater adoption of DLT on cloud computing platforms.

6. Conclusion

This report explains how Distributed Ledger Technology (DLT) would enhance the fault tolerance of cloud computing systems. Based on the experiments, the integration of DLT is seen to enhance the performance of cloud systems significantly through lessened fault detection and failure recovery delay, peak system up-time, and data integrity in fail operations. DLT's immutable and decentralised architecture facilitates real-time detection of faults together with faster recovery so that failure does not impair the integrity of the cloud infrastructure. This offers DLT as a suitable solution for cloud providers looking to increase the infrastructural reliability and resilience. While consolidation of DLT into cloud infrastructure is susceptible to extra energy consumption from the blockchain requirements of transaction processing, the return is in the guise of significantly enhanced fault tolerance, system stability, and data integrity.

The ability of DLT to maintain an open book of unaltered system events also bestows upon it a superior fault detection and recovery mechanism, offering a superior, secure and reliable platform than cloud computing. Future endeavours will nonetheless be necessitated to emphasise energy efficiency in the case of DLT systems, especially through less computationally costly consensus protocols. In addition, effective implementation of DLT on cloud infrastructure can further encompass factors such as usability, scalability, and long-term performance. Some of the problems these involve include those issues that have already been addressed by DLT and which can be used as a baseline in bringing cloud infrastructure towards higher stability, security, and efficiency in the provision of business and organisational processes that depend on cloud computing.

6.1. Constraints

Although Distributed Ledger Technology (DLT) has proven to be valuable when integrated with cloud infrastructure, there are various limitations to remember. The first of these is the extra computation burden of upholding a blockchain. Proof-of-work or proof-of-stake consensus algorithms employed by DLT systems consume massive computing resources. Such overhead would result in higher energy usage, especially in cloud-based deployments at scale with numerous nodes and transactions. Although this paper assumes the value of DLT exceeds its energy consumption, operational efficiency in blockchain remains an Issue. The second one is the scalability of DLT in the cloud environment. Scale-up in the cloud environment is equivalent to scaling up the blockchain ledger and transactions, which can result in slower transaction processing times and increased storage requirements.

Scalability of the blockchain system is an open issue, but mechanisms for successfully handling a large volume of transactions and data are yet to be developed. Also, the integration of DLT with existing cloud infrastructure will not be simple or cheap. Cloud providers will need to reconfigure their infrastructures to include the blockchain mechanism, which could involve significant restructuring, staff retraining, and addressing compatibility issues. These can be stumbling blocks to integrating DLT into cloud computing. Finally, DLT is relatively new to enterprise-scale implementation, and its real-world deployment in cloud systems remains to be extensively tested. Therefore, the practical reliability, cost-effectiveness, and security of DLT-based cloud systems can only be further empirically validated through widespread large-scale deployments.

6.2. Future Scope

The future for embracing Distributed Ledger Technology (DLT) in cloud-based systems is promising, with several areas for further research and development. Among the most important tasks is scaling DLT onto cloud infrastructure to the greatest extent possible. With higher levels of cloud infrastructure, transaction volumes get added to the blockchain. Implementations are also worked upon by developers, such as sharding, sidechains, and optimisation of the consensus algorithms (e.g., proof of stake, federated consensus), further to increase the scalability and efficiency of the DLT system. These developments can render DLT even more useful to gigantic cloud infrastructures without affecting performance or incurring a high price. Future research in yet another extremely crucial area is its coupling with emerging technologies like edge computing and 5G networks. As the cloud model evolves towards edge and decentralised computing systems, integrating it with blockchain and edge nodes can enhance the fault-tolerant capabilities of edge computers.

Apart from that, integration with IoT (Internet of Things) can also lead to new decentralised cloud system architectures where devices and nodes communicate directly with one another without applying any fault-tolerant approach. Though the energy consumption of DLT-based systems remains a matter of greatest concern, future studies need to focus on the development of energy-efficient blockchain algorithms. Hybrid deployment, combining conventional cloud fault tolerance methods with DLT, can also be explored for performance optimisation with moderate energy consumption. Finally, large-scale real-world deployments of DLT on cloud platforms would also substantiate the theoretical results of this research at large scales. Future research would involve collaborating with cloud firms to organise pilot projects, identifying fault tolerance strength, reducing costs, and cloud network sufficiency with the integration of DLT on a global scale.

Acknowledgement: The author gratefully acknowledges General Motors for their support and valuable resources that contributed to this work. Special thanks to their team for providing insights that enhanced the quality of the research.

Data Availability Statement: The study utilises a dataset related to Improving Fault Tolerance in Cloud-Based Systems Through Distributed Ledger Technology.

Funding Statement: This research received no financial support for its preparation or execution.

Conflicts of Interest Statement: The author declares no conflicts of interest. All sources have been properly cited in the manuscript.

Ethics and Consent Statement: The author obtained consent from the organization and all individual participants during data collection, with ethical approval and participant consent duly secured.

References

1. A. Rghioui, "Managing patient medical record using blockchain in developing countries: Challenges and security issues," in *2020 IEEE International conference of Moroccan Geomatics (Morgeo*), Casablanca, Morocco, 2020.
2. Y. Zhao, M. Cui, L. Zheng, R. Zhang, L. Meng, D. Gao, and Y. Zhang, "Research on electronic medical record access control based on blockchain," *Int. J. Distrib. Sens. Netw.*, vol. 15, no. 11, pp. 1–13, 2019.
3. H. Trollman, G. Garcia-Garcia, S. Jagtap, and F. Trollman, "Blockchain for ecologically embedded coffee supply chains," *Logistics*, vol. 6, no. 3, pp. 1-17, 2022.
4. S. F. Sayyad, M. Pawar, A. Patil, V. Pathare, and P. Poduval, "Features of blockchain voting: A survey," *Int. J. Innov. Res. Sci. Technol.*, vol. 5, no. 9, pp. 12–14, 2019.
5. X. Li, P. Jiang, T. Chen, X. Luo, and Q. Wen, "A survey on the security of blockchain systems," *Future Gener. Comput. Syst.*, vol. 107, no. 6, pp. 841–853, 2020.
6. Y. Guo and C. Liang, "Blockchain application and outlook in the banking industry," *Financ. Innov.*, vol. 2, no. 1, pp. 1-12, 2016.
7. S. Abdullah, J. Arshad, M. M. Khan, M. Alazab, and K. Salah, "PRISED tangle: a privacy-aware framework for smart healthcare data sharing using IOTA tangle," *Complex Intell. Syst.*, vol. 9, no. 3, pp. 3023–3041, 2023.
8. I.-C. Lin, C.-C. Chang, and Y.-S. Chang, "Data security and preservation mechanisms for industrial control network using IOTA," *Symmetry (Basel)*, vol. 14, no. 2, p. 237, 2022.
9. Y. Zhao, Z. Wang, J. Ni, X. Zhang, Y. Wang, and X. Lin, "Blockchain-based secure electronic health record system," *Int. J. Web Grid Serv.*, vol. 16, no. 1, pp. 15–30, 2020.
10. L. Luu, D. Chu, H. Olickel, P. Saxena, and A. Hobor, "Making smart contracts smarter," in Proc. *23rd ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)*, Vienna, Austria, 2016.