

Decent Cloud: A Decentralized Cloud Platform That You and Your Business Can Trust

Anonymous

August 25, 2024

Abstract

This paper introduces a novel decentralized cloud platform that bridges the gap between traditional centralized cloud services and fully decentralized systems. Our platform enables developers to access tailored virtual or physical servers, offered by identifiable node providers with established reputations and backed by robust legal contracts.

Unlike typical decentralized platforms, our approach requires node providers to build long-term reputations, allowing users to select providers based on their reliability and performance. This reputation-based system addresses key challenges in both centralized and existing decentralized cloud models, including vendor lock-in, opaque pricing, and barriers to entry.

In addition to the regular servers, our platform also supports sensitive data processing through Confidential Computing VMs and facilitates GPU node rental for ML and AI applications at competitive rates. By running the control plane on a blockchain while keeping the performance-sensitive data plane on the regular Internet, we overcome limitations of conventional web3 platforms, allowing seamless integration of legacy applications.

*Decent Cloud offers substantial benefits for both node providers and developers. **Developers** can (a) find cloud providers suitable for their specific needs, (b) obtain legal guarantees and SLA agreements typically expected by businesses, and (c) utilize services from multiple providers (multi-cloud deployments) with minimal added complexity. For **Node Providers**, it provides access to the trillion-dollar crypto market and this way reach a vast global user base. Crucially for node providers, Decent Cloud mitigates the "race to the bottom" pricing that plagues other decentralized cloud computing platforms, by implementing transparent, decentralized pricing and a reputation system that encourages and rewards honest behavior.*

1 Introduction

Cloud computing has revolutionized data storage and processing, offering on-demand access to vast computational resources without direct user management. However, the current landscape is dominated by a few major providers, leading to concerns about vendor lock-in, security vulnerabilities, and resource inefficiencies.

This paper presents a novel decentralized cloud platform that addresses these challenges while leveraging the strengths of both centralized and decentralized models. Our approach combines the reliability and performance of traditional cloud services with the flexibility and democratization of decentralized systems, creating a more robust, efficient, and user-centric cloud computing ecosystem.

1.1 Background and Motivation

In response to these challenges, decentralized cloud computing has emerged as a promising alternative. This approach distributes data and processes across multiple machines, or nodes, each equipped with its own processing power and storage. This distribution can lead to more efficient resource allocation and utilization, as resources can be dynamically shared among nodes based on demand.

Our objective is to address the inherent issues of

centralized cloud computing by constructing a secure and trustworthy system. This is achieved through the implementation of robust security measures and a reputation system, which collectively enhance the reliability and integrity of the platform.

Furthermore, we aim to promote transparency and democratic principles within the system by employing a decentralized matching engine and a Decentralized Autonomous Organization (DAO) model for governance. This approach is intended to stimulate competition and innovation in the field.

By lowering the barriers to becoming a cloud provider and offering a global rating system for all cloud providers, we enable developers to select providers based on their specific needs and preferences. This approach not only fosters a more competitive and innovative environment but also empowers developers by providing them with more choice and control over their cloud computing solutions.

1.2 Academic research

Decentralized cloud computing has garnered significant attention in the academic community as a viable alternative to centralized cloud models. Research in this domain explores various approaches to enhancing the efficiency, security, and scalability of cloud services through decentralization.

Xu Chen's work on "Decentralized Computation Offloading Game for Mobile Cloud Computing" [1]

proposes a game-theoretic framework for efficient computation offloading in mobile cloud environments. This work highlights the potential of decentralized models in optimizing resource allocation, particularly in scenarios involving mobile devices with limited computational capabilities. Similarly, Sladana Jošilo and G. Dán in their paper "Selfish Decentralized Computation Offloading for Mobile Cloud Computing in Dense Wireless Networks" [2] investigate the dynamics of decentralized computation offloading in dense wireless networks. Their research underscores the challenges of coordinating decentralized nodes and the importance of incentive mechanisms in ensuring efficient resource sharing.

Kunal et al. [3] explore fog computing as a decentralized layer between users and cloud servers, aiming to reduce latency and enhance service delivery. This research aligns with the broader goal of decentralization by distributing computational tasks closer to the edge, thereby improving responsiveness and reducing the load on central servers.

Nguyen et al. [4] provide a comprehensive review of integrating blockchain with Cloud of Things (CoT), emphasizing the potential of blockchain to address decentralization challenges, enhance data privacy, and secure network transactions. Their work provides insights into the benefits of combining blockchain technology with decentralized cloud architectures.

Wenjuan Li et al. proposed a blockchain-based trust management system in cloud computing [5], which aligns closely with our platform's reputation-based approach. Their research illustrates how blockchain can be leveraged to establish trust among decentralized nodes, ensuring reliable and secure service delivery.

Chegini et al. [6] discuss the automation of tasks and processes in the IoT-Fog-Cloud ecosystem, addressing big data and heterogeneity challenges. Their work highlights the need for decentralized solutions to manage the complexity and scale of modern cloud environments effectively.

Finally, Espinel Sarmiento et al. [7] survey the application of Software Defined Networks (SDN) in managing distributed Cloud-Edge infrastructures. This research is particularly relevant to our platform, as it demonstrates how SDN can facilitate the coordination and management of decentralized resources across a distributed cloud environment.

1.3 Industry projects

In recent years, several industry projects have attempted to build decentralized cloud platforms. These platforms leverage blockchain technology to create a network of independent node providers that

can store data or run arbitrary workloads in a decentralized fashion, and highly resilient to malicious actors.

- **Filecoin** (<https://filecoin.io/>): Filecoin is a decentralized storage network designed to store and retrieve data in a distributed fashion. By incentivizing storage providers through its native cryptocurrency, Filecoin creates a competitive marketplace for data storage. While Filecoin excels in decentralized storage, it does not extend to compute resources. In contrast, our platform aims to integrate both storage and compute capabilities, offering a more holistic cloud solution.
- **Storj** (<https://www.storj.io/>): Storj is a decentralized cloud storage platform that focuses on secure and distributed data storage. It fragments and encrypts data before distributing it across its network, ensuring high levels of security and redundancy. Similar to Filecoin, Storj is storage-centric and does not offer compute resources. Our platform distinguishes itself by providing a marketplace that includes both storage and compute services.
- **Sia** (<https://sia.tech/>): Sia operates as a decentralized storage platform that leverages global hard drive capacity to create a distributed data storage network. While Sia effectively reduces storage costs and enhances data redundancy, it lacks the ability to offer compute resources. Additionally, it does not incorporate a reputation-based system for node providers. Our platform addresses these gaps by integrating compute capabilities and establishing a reputation system to enhance trust and reliability.
- **Arweave** (<https://www.arweave.org/>): Arweave offers a unique approach to data storage by providing permanent and sustainable storage solutions. It is particularly suited for archiving and storing immutable data. However, Arweave's focus on storage limits its applicability in scenarios requiring compute resources. Our platform expands on this by supporting both storage and compute, catering to a broader range of cloud computing needs.
- **DFINITY** (<https://dfinity.org/>): DFINITY's Internet Computer represents a significant step toward decentralized cloud computing by enabling developers to host entire applications on its decentralized network. Despite its innovation, the Internet Computer's reliance on the Consensus protocol imposes performance limitations, particularly for high-throughput applications. Our platform, by contrast, separates the control and data planes, allowing for higher performance and broader applicability across

diverse workloads.

- **Akash Network** (<https://akash.network/>): Akash Network is a decentralized cloud computing platform that focuses on offering compute resources through a permissionless marketplace. Akash allows users to rent computing power from providers across the globe, targeting workloads such as web hosting, machine learning, and containerized applications. While Akash shares similarities with our platform, our solution emphasizes a more comprehensive approach, integrating robust reputation systems and legal frameworks to ensure quality and reliability, which many users see as essential.
- **Golem** (<https://golem.network/>): Golem functions as a decentralized supercomputer that aggregates the computational power of user machines worldwide. It is particularly well-suited for compute-intensive tasks such as rendering and scientific simulations. However, Golem's focus on specialized compute tasks limits its versatility as a general-purpose cloud platform. Our platform aims to provide a more flexible solution that accommodates a wide range of workloads, from basic web hosting to complex AI processing.
- **iExec** (<https://iex.ec/>): iExec provides a decentralized marketplace for cloud computing resources, primarily serving decentralized applications (DApps) on blockchain networks. While iExec offers an efficient solution for DApp developers, its reliance on the Consensus protocol can slow down execution times, limiting its use for performance-sensitive applications. Our platform, by contrast, is designed to serve a broader spectrum of applications, with a focus on minimizing latency and maximizing performance.
- **Aethir** (<https://www.myedge.io/>): Aethir offers a decentralized marketplace for renting GPU resources, specifically targeting machine learning (ML) and artificial intelligence (AI) workloads. By focusing on GPU availability, Aethir addresses a critical need in the AI and ML communities. Our platform incorporates similar GPU rental capabilities but extends the offering to include other compute and storage resources, providing a more comprehensive cloud solution.
- **Cudos** (<https://www.cudos.org/>): Cudos is a decentralized cloud computing platform that aims to bridge the gap between cloud and blockchain by providing compute resources for blockchain and non-blockchain applications. Cudos supports a wide range of workloads, from simple hosting to complex decentralized finance (DeFi) applications. Our platform builds on

these ideas by offering an even more integrated approach, combining compute, storage, and reputation systems within a single ecosystem.

- **ICN (Impossible Cloud Network)** (<https://icn.global/>): The Impossible Cloud Network (ICN) is a decentralized approach to cloud infrastructure, providing a platform for managing hardware and cloud resources. It leverages decentralized technologies to ensure data security, integrity, and resiliency. ICN's approach to managing global cloud resources aligns with our platform's goal of mitigating vendor lock-in while enhancing transparency and control. Our platform further differentiates itself by integrating a reputation system and legal frameworks to ensure service quality and compliance.
- **Fluence** (<https://fluence.dev/>): Fluence offers a decentralized serverless platform that enables developers to deploy applications without the need for centralized servers. It provides a peer-to-peer network for executing code, using a special-purpose language called Aqua for cryptographic proofs. While Fluence's approach is innovative, the need for specialized language limits its flexibility. Our platform, in contrast, supports a broader range of programming languages and environments, and integrates a reputation system to ensure service quality and compliance, making it accessible to a wider audience and a wider range of use cases.

These works highlight the growing interest in decentralized cloud computing and the potential benefits it can offer. However, despite these efforts there is still a need for a comprehensive solution that addresses the issues of the current cloud computing paradigm while leveraging the benefits of decentralization. This is the problem that our proposed decentralized cloud platform aims to solve and in the following sections we provide an outline of how we plan to do this.

2 Problem Statement

The current cloud computing landscape is dominated by a handful of large providers, including Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure, which together control roughly 70% of the cloud computing market, as of 2024. While these providers offer a comprehensive range of services spanning basic compute and storage to advanced machine learning and analytics tools, this centralized model presents several significant challenges:

- **Vendor lock-in:** The choice of a cloud provider




Figure 1: An illustration of the proposed solution.

often becomes a long-term commitment due to:

- Proprietary APIs and services that are not easily portable
- Substantial costs and technical challenges associated with migrating data and applications
- Pricing models that incentivize increased usage within a single ecosystem
- **Limited control and transparency:** Users face significant constraints in managing their cloud resources:
 - Lack of granular control over data storage locations and processing methods
 - Limited visibility into the underlying infrastructure and its performance characteristics
 - Difficulty in optimizing applications due to the abstracted nature of cloud services
- **Data privacy and security concerns:** Centralized cloud providers present attractive targets for cyberattacks:
 - A single breach can potentially compromise vast amounts of data across multiple clients
 - Users must trust the provider’s security measures without full transparency
 - Compliance with varying international data protection regulations becomes complex
- **Cost and resource inefficiencies:** The pricing structures of major cloud providers can lead to financial inefficiencies:
 - Complex, often opaque pricing models make cost prediction and management challenging
 - Users frequently over-provision resources to ensure performance, leading to waste
 - Lack of a true pay-per-use model for many services results in unnecessary expenses

These issues have varying impacts across different stakeholders in the cloud computing ecosystem:

- **Individual Developers:** Face barriers to experimentation and innovation due to high costs and complex management of cloud resources. They may struggle to scale their projects efficiently as requirements grow.
- **Startups:** Often find themselves locked into a specific provider early in their development, limiting future flexibility. They may face sudden

cost increases as they scale, impacting their financial planning and viability.

- **Enterprises:** Struggle with vendor lock-in, which can impact long-term strategic decisions. They face significant challenges in maintaining regulatory compliance across different regions and in protecting sensitive data.
- **Government and Non-Profit Organizations:** May face difficulties in ensuring data sovereignty and meeting strict security requirements. Budget constraints can limit their ability to leverage advanced cloud services effectively.

Addressing these challenges is crucial for fostering a more open, efficient, and innovative cloud computing ecosystem. The need for a solution that provides greater flexibility, transparency, and cost-effectiveness while maintaining robust security and performance is evident across all sectors of the industry.

3 Proposed Solution

To address the limitations of the current centralized cloud computing model, we propose a novel decentralized cloud platform that combines the best aspects of traditional cloud services with the advantages of decentralized systems. Our solution is designed to tackle each of the key issues identified in the problem statement while introducing innovative features that enhance flexibility, security, and efficiency.

3.1 Core Components of the Solution

Decentralized Infrastructure Our platform leverages a network of independent node providers, each capable of hosting and executing arbitrary workloads. This approach:

- Mitigates vendor lock-in by offering a diverse ecosystem of providers
- Enhances resilience by distributing resources across nodes of different providers
- Improves resource utilization by tapping into underused computational capacity

Reputation-Based Trust System Unlike typical Web3 solutions that rely on consensus mechanisms to mitigate the Byzantine Fault Tolerance (BFT) problem

— allowing up to 50% (or a bit less on some platforms) of the network to be malicious — our platform takes a different approach. We primarily rely on conventional legal contracts and Service Level Agreements (SLAs), supplemented by a robust blockchain-tracked reputation system. This dual approach ensures accountability through traditional means while providing an additional layer of trust and transparency.

In the anonymous blockchain space, an actor is more likely to behave maliciously than in the real world. In the real world there is a lot more at stake than in the blockchain space, an actor would be much less likely to behave maliciously, given that they risk suffering both legal and financial consequences for malicious behavior. Malicious behavior may still occur, but the potential benefit of such behavior needs to clearly compensate the potential cost, which may come in the form of a loss of the long-term financial benefit and the potential financial loss in the following legal processes, as analyzed in more details in Section 4.

Decent Cloud is designed with this in mind, and it focuses on incentivizing honest behavior and providing high-quality service while implementing mechanisms to penalize poor performance or malicious actions. Good reputation is hard to earn, and it's easy to lose. Especially if big clients (with high reputation) are dissatisfied with the provided services.

Building Reputation: Node providers and developers accrue reputation through successful transactions and positive interactions on the platform:

- Each completed transaction between a developer and a node provider results in a 2% transaction fee, from which there is a 1% increase in both parties' reputation scores.
- Consistent user satisfaction contributes to gradual reputation increases of node providers over extended periods and numerous transactions.
- Developers who rent many resources or spend significant amounts over longer periods of time will have higher reputation scores than those who do not.

Reputation Dynamics: Our system introduces an innovative approach to managing reputation:

- Both node providers and developers build reputation through successful transactions.
- Rather than complicating the platform to autonomously track and judge honest behavior, we empower users to report and penalize malicious behavior.
- If dissatisfied with a service or collaboration, a user can choose to reduce another party's reputation. This action comes at a cost: the sender's

reputation decreases by the amount spent, while the receiver's reputation decreases by a higher amount, such as 200

This mechanism offers several advantages:

- **Robustness:** It's straightforward to implement and reason about.
- **Quality Incentive:** It motivates node providers to maintain high-quality service and prioritize user satisfaction.
- **Resilience against Malicious Actors:** Malicious node providers would quickly acquire poor reputations, limiting their ability to attract new users. Similarly, malicious developers quickly lose power to hurt the reputation of others, as they lose their own reputation in the process.

Reputation System Impact: This dynamic reputation system empowers users to make informed decisions when selecting node providers. It addresses trust and security concerns associated with decentralized systems while preserving the performance advantages of traditional cloud services. All reputation scores are publicly visible and immutably recorded on the blockchain, ensuring transparency and creating a reliable history of interactions and service quality.

Efficient Resource Allocation Our platform can be used to ensure optimal resource allocation. For instance:

- Task requirements can be matched with available node resources in real-time
- Dynamic load balancing optimizes resource utilization across the network
- Pricing is transparent and competitive, driven by market demand

This approach allows users to tackle the cost inefficiencies and opaque pricing models of traditional centralized providers.

Enhanced Security and Privacy The system supports sensitive data processing and storage using Confidential Computing VMs, and also facilitates the rental of GPU nodes for Machine Learning (ML) and Artificial Intelligence (AI) training and inference applications at reasonable market prices, catering to different developer preferences. Some developers may opt for the lowest cost GPUs irrespective of node provider reputation and confidentiality guarantees, while others may prefer high-end GPUs with higher node provider reputation. This flexibility accommodates all developer categories, a feature unique to this platform.

For some use cases, such as storing and rotating tokens, certificates, and other secrets, developers will need nodes with additional security. For these use cases, platform will support Confidential Containers[8].

To address data privacy and security concerns, our platform offers:

- Support for Confidential Computing VMs, enabling secure processing of sensitive data
- End-to-end encryption for data in transit and at rest
- Granular control over data location and processing parameters

3.2 Addressing Specific Challenges

Overcoming Vendor Lock-in Our platform’s open architecture and standardized interfaces allow for easy migration between providers, addressing the vendor lock-in issue prevalent in centralized systems.

Improving Control and Transparency Users have unprecedented visibility into the underlying infrastructure and can choose specific nodes or node providers based on their requirements, enhancing control and transparency.

Enhancing Cost-Effectiveness The platform’s market-driven pricing model and efficient resource allocation system ensure that users only pay for the resources they actually use, addressing the cost inefficiencies of traditional cloud services.

3.3 Innovative Features

Wide Range of Applications Our platform supports diverse applications, from scientific high-performance computing and machine learning to web hosting and data storage, catering to a broad spectrum of user needs.

DAO Governance We implement a Decentralized Autonomous Organization (DAO) model for platform governance, ensuring transparency and community-driven decision-making.

Blockchain-Based Control Plane By running the control plane on a blockchain, we ensure transparent and immutable record-keeping for financial transactions and reputation tracking, while keeping the performance-sensitive data plane on the regular internet for optimal speed.

3.4 Benefits for Stakeholders

For Developers

- Access to a diverse range of computational resources

- Flexibility to choose providers based on specific needs
- Potential for significant cost savings

For Node Providers

- Opportunity to monetize unused computational resources
- Access to a global market of users
- Incentives for maintaining high-quality service through the reputation system

Figure 2 illustrates the architecture of our proposed platform, showcasing how these components interact to create a robust, efficient, and user-centric cloud computing ecosystem.

In the following sections, we will delve deeper into the technical details of each component, exploring how they work together to realize the full potential of decentralized cloud computing.

4 Incentive Analysis for Malicious Behavior in Decentralized Systems

4.1 Introduction

Understanding the incentives that could drive malicious behavior is crucial for enhancing the security and robustness of decentralized systems. The aim is to provide a nuanced understanding of the conditions under which malicious activities are more likely, offering insights for designing secure systems that do not need to rely on conventional consensus-based BFT schemes.

4.2 Model Variables

- $P(B_m), P(B_h)$: Probability distributions for the expected benefits from malicious and honest behavior.
- $P(C_m), P(C_h)$: Probability distributions for the financial costs.
- $P(R_m), P(R_h)$: Probability distributions for the financial equivalent of reputation loss or gain, respectively for malicious and honest behavior.
- $P(T_m), P(T_h)$: Probability distributions for the time required for malicious and honest behavior.

4.3 Model Formulation

The expected net benefits, adjusted for time, can be modeled as:

$$E[N_m] = \frac{E[B_m]}{E[T_m]} - (E[C_m] + E[R_m])$$

$$E[N_h] = \frac{E[B_h]}{E[T_h]} - (E[C_h] + E[R_h])$$

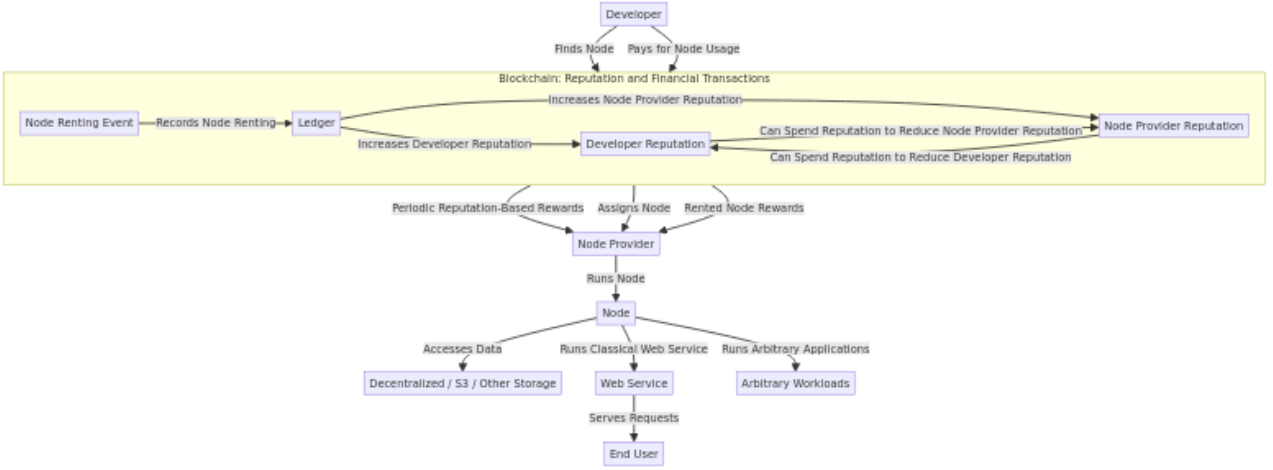


Figure 2: Proposed platform architecture.

Here, $E[C_m] + E[R_m]$ and $E[C_h] + E[R_h]$ represent the total expected costs for malicious and honest behavior, respectively. These costs include both the financial costs and the financial equivalent of reputation loss.

4.4 Conditions for Malicious Behavior

Malicious behavior is more likely if $E[N_m] > E[N_h]$.

4.5 Proof

Theorem: Malicious behavior is more likely to occur if and only if $E[N_m] > E[N_h]$.

Proof:

We aim to prove that malicious behavior is more likely when the expected net benefit of malicious behavior $E[N_m]$ is greater than that of honest behavior $E[N_h]$.

The condition $E[N_m] > E[N_h]$ simplifies to:

$$\frac{E[B_m]}{E[T_m]} - (E[C_m] + E[R_m]) > \frac{E[B_h]}{E[T_h]} - (E[C_h] + E[R_h])$$

Rearranging terms, we get:

$$\frac{E[B_m]}{E[T_m]} - \frac{E[B_h]}{E[T_h]} > (E[C_h] + E[R_h]) - (E[C_m] + E[R_m])$$

This inequality shows that malicious behavior is more likely when the expected time-adjusted benefit of malicious behavior over honest behavior is greater than the expected net cost difference, which includes both the financial costs and the financial equivalent of reputation loss for both behaviors.

Therefore, malicious behavior is more likely to occur if and only if $E[N_m] > E[N_h]$.

Q.E.D.

5 Implementation

The proposed decentralized cloud platform consists of several key components:

- **Nodes and Networking:** The platform runs numerous nodes, and each node adds its resources such as processing power, storage, and bandwidth. The network functions on a divide-and-conquer principle, minimizing communication between different applications running on the network. This strategy enables the network to scale and accommodate numerous nodes and tasks.
- **Reputation System:** The platform uses a blockchain-based system to ensure the security, transparency, and fairness of resource allocation. Smart contracts are utilized to verify that node providers, nodes, and developers are all acting appropriately, and that each party is suitably rewarded for their contributions or penalized for any misconduct.
- **Token Economics:** The platform employs a distinct dual-token system: one that can be used for transactional purposes and another for reputation tracking. The transactional token facilitates financial interactions between node providers and developers, while the reputation token, non-exchangeable in nature, is earned by node providers and developers during transactions.
- **Application Support:** The platform supports a wide range of applications, from scientific computing and machine learning to web hosting and data storage. The platform provides APIs and SDKs that simplify the process for developers to build and deploy their applications on the platform.

5.1 Nodes and Networking

The foundation of the proposed decentralized cloud platform lies in its nodes and the networking infrastructure that connects them. The platform operates on a multitude of nodes, each contributing resources such as processing power, storage, and bandwidth. This section provides a detailed overview of the nodes and networking component of the platform.

Nodes In the context of our platform, a node is a computing device that participates in the network. Each node contributes its computational resources, including processing power (CPU), memory (RAM), storage (hard disk or SSD), and network bandwidth. Nodes can vary in size, from personal computers to large data center servers. This diversity in node capabilities is a strength of the platform, as it allows for a wide range of application requirements.

Each node operates a software client that enables its participation in the network. This client manages the node's resources, communicates with other nodes, and executes assigned tasks. The client software is designed to be lightweight and efficient, minimizing its impact on the node's performance.

Networking The networking component of the platform connects nodes and facilitates their communication. The network employs a divide-and-conquer approach, where nodes belonging to a specific developer form a smaller subnetwork or cluster. These clusters maintain direct communication within themselves, and communication between different clusters is minimal unless explicitly established by the developer.

Node providers, who typically host many Nodes or Virtual Machines (VMs), do not need to provide a public IPv4 address for each Node. This is crucial as IPv4 addresses are a scarce resource. Instead, node providers only need to port forward, thereby publicly exposing a few ports for each Node. These ports can be used for peering different nodes of the same developer. For example, some developers may peer nodes with Nats.io, while others might use WireGuard, but any other solution of the developer's choice would work.

Node providers also need to have HTTP and HTTPS ingress, on which a reverse proxy is set up. This reverse proxy forwards requests to the appropriate VM using Server Name Indication (SNI) technology, i.e., based on the target domain name. This routing can be done *without* SSL termination to prevent Node providers from being able to snoop all ingress traffic. For this, tools like HAProxy or recent versions of Nginx can be used. This approach improves security, as it prevents node providers from inspecting traffic to and from nodes.

The illustrated networking approach minimizes communication between nodes in different clusters, reducing network congestion and improving performance.

Security While premium security is a necessity, it often comes with a cost. Techniques such as homomorphic encryption, although secure, can significantly slow down operations and inflate data size by a factor of 10 to 100 or more.

Technologies like Intel SGX and TDX, AMD SEV-SNP, ARM TrustZone, and open-source efforts like UC Berkeley's Keystone offer robust security at a relatively low cost. However, the use of Confidential Computing VMs can lead to a slight performance loss and added complexity. Certain operations are currently unsupported with Confidential Computing VMs:

1. Nested virtualization: Running a VM within another VM is not feasible.
2. Memory overcommitment: Confidential VMs' encrypted memory limits the ability to share or swap out memory, hindering memory overcommitment.
3. VM snapshots, cloning, and live migration: The hardware-specific encryption keys make it impossible to migrate the encrypted memory state to another host, increasing downtime risk.
4. Debugging and Inspection: The VM's encrypted and isolated memory state complicates troubleshooting and monitoring.
5. High-performance Computing (HPC): Direct, low-latency access to hardware resources required by some HPC workloads can be complicated by the additional encryption and isolation layer.
6. Hardware Compatibility: Not all hardware platforms or devices are compatible with confidential computing technologies.

Despite these challenges, Confidential Computing VMs are invaluable in scenarios where sensitive data needs to be processed or stored securely. These include:

1. Private Token Management and Rotation: Confidential computing securely handles sensitive operations like token management and rotation.
2. SSL Certificate Renewal and Key Storage: Renewing SSL certificates and storing private keys can be securely managed within a confidential computing environment.
3. Secure Multi-party Computation: Confidential computing provides a secure environment for computation in scenarios where multiple parties need to collaborate on sensitive data.




Figure 3: Networking Diagram. Note that blockchain (and the consensus protocol) are not on the data path, so the node performance isn't limited by the consensus speed. Blockchain is on the control path only and is primarily used for tracking reputation.

4. Data Privacy and Compliance: Confidential computing can help meet strict data privacy and compliance requirements in industries handling sensitive data.
5. Intellectual Property Protection: Confidential computing ensures that sensitive intellectual property, such as proprietary algorithms or models, are not exposed.

The proposed strategy is to facilitate the use of both regular and Confidential Computing VMs, thereby providing developers with the flexibility to select VM types based on their specific requirements. For example, developers could manage confidential workloads using Confidential Computing VMs, while employing regular VMs for all other workloads. This approach promotes a tailored and efficient use of resources, aligning with the individual needs of each workload.

5.2 Reputation System

The reputation system is the cornerstone of our platform, maintaining simplicity comparable to existing cloud platforms while reliably running conventional workloads on a globally distributed network of node providers. It enhances security, transparency, and fairness in resource allocation. All participants' reputations, including developers and node providers, are tracked via smart contracts on a blockchain.

When a developer pays a node provider, the system charges the node provider a 2% fee. A portion of this fee is automatically transferred to a DAO-controlled wallet that finances research and development efforts. In return, the reputations of both the developer and the node provider increase by 1%.

Assuming mostly honest behavior, the reputations of developers and node providers will gradually grow, proportional to the service payment amounts and transactions performed on the platform. All reputations are publicly visible, allowing, but not forcing, node providers and developers to restrict their interactions to high-reputation parties. Additionally, node providers' reputations are used to reward them for offering nodes for rent, making it crucial for all node providers to maintain a high reputation.

If either a node provider or a developer is dissatisfied with the service or collaboration, they can spend their reputation to reduce the other's reputation. The

sender's reputation decreases by the amount spent, while the receiver's reputation decreases by a higher amount, such as 200%, to incentivize honest and cooperative behavior. For members with high reputations, these reductions can be severe in extreme cases. Spending hard-earned reputation to reduce another party's reputation is not undertaken lightly, as rebuilding reputation takes significant time and money.

Reputation is hard to earn and easy to lose, encouraging all parties to uphold high standards of quality, communication, and integrity at all times, while still offering a means to penalize potentially dishonest or malicious behavior, in a general and transparent manner.

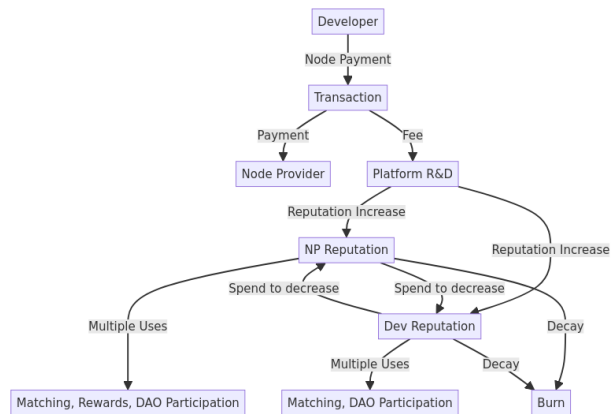


Figure 4: Reputation tracking system, running on Blockchain.

5.3 Resource Allocation

The platform employs a matching engine and, if necessary, machine learning algorithms to efficiently allocate resources and match developers with suitable nodes.

A developer specifies their requirements, which may include some or all of the following:

- Acceptable price
- Required virtual CPU (vCPU)
- Required Random Access Memory (RAM)
- Required storage (which can be unlimited)
- Required network bandwidth (which can be unlimited)
- Minimum node provider reputation
- Renting period

The matching engine processes these requirements and returns a list of suitable nodes. The developer then selects one or more nodes from this list and sends a request to confirm the rental. Upon receiving payment from the developer, the matching engine provides a rental contract ID. The developer can use this contract ID to start utilizing the node.

If the developer wishes to extend the rental period, they can make another request at a later time. The process for extending the rental period is similar to the original rental process, with the key difference being that the developer already has access to the node.

5.4 Token Economics (Tokenomics)

The platform tokenomics balances two key components: (1) the Decentralized Cloud Token (DCT), with which all key platform operations are paid, and (2) the actor reputation, which is hard to earn and easy to lose. The reputation is discussed in more details in Section 5.2.

Token Minting (Generation) and Distribution. Similar to Bitcoin, DC tokens are minted in each new block, occurring approximately every 10 minutes. The initial block mints 50 tokens, and after that the minted tokens are halving every 210,000 blocks. This model is designed to sustain token supply growth at a decreasing rate, with a final cap due to the smallest unit set by the DC token's nine decimal places, calculated to total of approximately 21,000,000 DC tokens.

The distribution of DC tokens received in a block are weighted by their reputation within the protocol. This reputation is dynamically adjusted based on participants' transaction activities and behaviors. The reward distribution uses a square root scaling function to calculate weights, providing diminishing returns on increased reputation scores, thereby maintaining incentives for both high and low-reputation participants.

Importantly, there will be no other DCT generation other than through this minting and rewarding process. There will be no additional token sale, airdrop, or any other form of token distribution. This approach promotes fair distribution of DCT across participants and balances supply and demand over time.

Project Funding. Participation in the DC token reward distribution requires a fee, set at 1/100th of the block reward, creating a steady funding stream for project-related activities. These fees are directed to predetermined wallets, which include known wallets designated for project development. This ensures transparency in funding flows and supports ongoing development efforts.

Stability, Governance, and Community Engagement. The stability and security of the token value are ensured through several mechanisms. Firstly, the supply of DCT is both limited and predictable, which, when balanced against the inherent demand created by developers needing to rent nodes, aids in maintaining market equilibrium. Secondly, the platform's Decentralized Autonomous Organization (DAO) has the capability to implement necessary adjustments to the reward system. This flexibility allows for the mitigation of extreme price volatility. Finally, the platform is committed to adhering to all relevant legal and regulatory standards, further ensuring the security of transactions and the integrity of the token system.

This model would facilitate community input through a transparent proposal system where community members can submit and vote on changes. It will also be possible to implement a delegated governance model, to address potential voter apathy and ensure informed decision-making.

Token Demand and Supply. Developers and node providers need to obtain DCT in order to pay for the registration fee. The DCT token can be obtained on an exchange, and the exchange obtains the token from one of the existing token holders. This creates an inherent demand for DCT. Node providers, on the other hand, will need to cover costs such as electricity and internet connectivity, which may necessitate selling a portion of their DCT. However, some providers may choose to retain a portion of their DCT in anticipation of a price increase, which could, in turn, increase the likelihood of such a price increase and creating a virtuous circle.

Smart Contract Interoperability. The system would run on the Internet Computer platform and DCT would adhere to the ICRC-1 and potentially also newer token standard to ensure compatibility and ease of integration with wallets and exchanges. In the future, an exploration of cross-chain functionality will take place, to enhance liquidity and utility across different blockchain ecosystems.

5.5 Application Support

The ability to support a wide range of applications is a critical factor for the success of any computing platform. This versatility not only broadens the platform's potential user base but also increases its utility and value for each user. By accommodating various types of applications, from scientific computing and machine learning to web hosting and data storage, the platform can cater to diverse needs and use cases, thereby attracting a wider audience and fostering a vibrant and diverse user community [9].

Our platform is designed with this versatility in

mind. It provides comprehensive support for a wide range of applications, enabling developers to leverage the platform's resources and capabilities to build and deploy various types of applications. This is in line with the best practices in platform design, which emphasize the importance of flexibility and extensibility in supporting diverse applications [10].

To facilitate application development and deployment, the platform provides Application Programming Interfaces (APIs) and Software Development Kits (SDKs). These tools simplify the process for developers to interact with the platform, abstracting away the complexities of the underlying infrastructure. This aligns with the principles of developer experience, which highlight the importance of providing developers with effective tools and resources to enhance their productivity and satisfaction [11].

Moreover, the platform is designed to be permissionless, meaning that the community of node providers and developers themselves can decide what kind of support should be added. For instance, there is nothing stopping node providers from adding nodes with GPU support, even though this is not something that is explicitly supported or required by the platform. This permissionless nature fosters innovation and allows the platform to adapt to the evolving needs of its community.

In conclusion, by providing comprehensive application support, continuously adapting to the evolving needs of developers, and allowing the community to drive innovation and expansion, our platform aims to be a versatile and valuable tool for a wide range of applications.

6 Benefits and Impact

The proposed decentralized cloud platform offers numerous potential benefits and could significantly impact the cloud computing industry:

- **Democratization of Resources:** By leveraging unused computational resources globally, the platform democratizes access to cloud computing. This could make cloud services more affordable and accessible, particularly in developing countries or remote areas.

Quantitative Projection: We estimate that this democratization could potentially reduce entry-level cloud computing costs by 30-50% compared to current centralized providers.

- **Cost Reduction:** The platform could significantly reduce cloud computing costs by eliminating the need for expensive data centers and using more efficient resource allocation.

Quantitative Projection: Based on our models, enterprises could see a 20-40% reduction in their

cloud computing expenses, potentially saving millions for large-scale operations.

- **Increased Competition:** By providing a decentralized alternative to traditional cloud providers, the platform promotes competition in the cloud computing market. This could drive down prices compared to conventional cloud providers and spur innovation.

Quantitative Estimate: We project a 15-25% increase in the number of viable cloud service providers within the first three years of platform adoption.

- **Environmental Sustainability:** Using existing computational resources more efficiently could reduce the need for new data centers, contributing to environmental sustainability.

Quantitative Projection: We estimate that widespread adoption of our platform could reduce the carbon footprint associated with cloud computing by 10-20% over five years.

- **Resilience and Security:** The decentralized nature of the platform enhances resilience against attacks and failures. The use of blockchain technology and our reputation system further bolsters security and transparency.

Quantitative Estimate: We project a 40-60% reduction in the impact of individual node failures or security breaches compared to centralized systems.

The impact of the platform could extend far beyond the cloud computing industry. By making cloud computing more affordable and accessible, it could enable new applications and services across various sectors:

- **Scientific Research:** Affordable access to high-performance computing could accelerate research in fields like genomics, climate modeling, and particle physics.

Case Study: A small university research team could leverage the platform to run complex climate models at 1/3 the cost of traditional cloud services, enabling more comprehensive studies on a limited budget.

- **Education:** Educational institutions could provide better computational resources to students, enhancing learning experiences in computer science, data analysis, and other tech-intensive fields.

Scenario: A rural high school could set up an advanced computing lab using our platform, giving students access to resources previously available only in well-funded urban schools.

- **Healthcare:** Improved access to computational resources could enhance medical imaging analy-

sis, drug discovery processes, and personalized medicine initiatives.

Case Study: A startup developing AI for early cancer detection could use our platform to process millions of medical images at 50% lower cost, potentially bringing their life-saving technology to market faster.

- **Small Business and Startups:** Reduced costs and easier access to cloud resources could lower barriers to entry for tech startups and enable small businesses to leverage advanced analytics and AI.

Scenario: A local retail chain could implement sophisticated inventory management and customer behavior analysis, previously only feasible for large corporations, improving their competitiveness.

In summary, our decentralized cloud platform has the potential to democratize access to advanced computing resources, foster innovation across multiple sectors, and contribute to a more sustainable and equitable digital economy. The projected cost savings, improved accessibility, and enhanced performance could catalyze breakthroughs in science, education, healthcare, and business, driving societal progress and economic growth.

7 Conclusion

In conclusion, the proposed decentralized cloud platform represents a significant leap forward in cloud computing technology. By leveraging distributed systems and blockchain technology, our platform addresses key limitations of both centralized and existing decentralized cloud solutions:

- It democratizes access to computational resources, reducing costs and stimulating competition.
- It enhances security and privacy through a robust reputation system and support for Confidential Computing.
- It improves resource utilization and environmental sustainability through efficient matching of supply and demand even across different node providers.
- It provides a seamless transition path for legacy applications while offering the benefits of decentralized systems.

The implementation of such a transformative platform undoubtedly presents challenges, particularly in areas of scalability, regulatory compliance, and user adoption. However, the potential benefits far outweigh these obstacles. Our platform is not merely

an incremental improvement to cloud computing; it represents a paradigm shift with far-reaching implications for the technology industry and society at large.

Unlike many preceding web3 projects, our platform does not require a complete overhaul of existing applications. Instead, it offers a bridge between traditional and decentralized cloud computing, allowing applications to operate as usual but on a more transparent, fair, and efficient infrastructure. This is the cloud as it should have been: resistant to vendor lock-in, responsive to user needs, and resilient against misconduct.

Looking ahead, there are numerous exciting avenues for future research and development:

- Exploring advanced machine learning techniques for even more efficient resource allocation and predictive scaling.
- Investigating integration with emerging technologies such as edge computing and 5G networks.
- Developing industry-specific solutions that leverage the platform's unique capabilities in fields like healthcare, finance, and scientific research.

We call upon the academic community, industry professionals, and policy makers to engage with this technology, to critically examine its potential, and to contribute to its development. The decentralized cloud is not just a technological innovation; it's a step towards a more open, efficient, and equitable digital infrastructure for all.

As we move forward, let us embrace this opportunity to reshape cloud computing and, by extension, the very fabric of our digital world. The potential for positive impact is immense, and the time to act is now.

References

- [1] Xu Chen. "Decentralized computation offloading game for mobile cloud computing". In: *IEEE Transactions on Parallel and Distributed Systems* 26.4 (2014), pp. 974–983.
- [2] Sladana Jošilo and G. Dán. "Selfish Decentralized Computation Offloading for Mobile Cloud Computing in Dense Wireless Networks". In: *IEEE Transactions on Mobile Computing* (2018). DOI: 10.1109/TMC.2018.2829874. URL: <https://dx.doi.org/10.1109/TMC.2018.2829874>.
- [3] Sourav Kunal, Arijit Saha, and Ruhul Amin. "An overview of cloud-fog computing: Architectures, applications with security challenges". In: *Security and Privacy* 2.4 (2019), e72.

- [4] Dinh C Nguyen et al. "Integration of blockchain and cloud of things: Architecture, applications and challenges". In: *IEEE Communications surveys & tutorials* 22.4 (2020), pp. 2521–2549.
- [5] Wenjuan Li et al. "Blockchain-based trust management in cloud computing systems: a taxonomy, review and future directions". In: *Journal of Cloud Computing* (2021). DOI: 10.1186/s13677-021-00247-5. URL: <https://dx.doi.org/10.1186/s13677-021-00247-5>.
- [6] Hossein Chegini et al. "Process automation in an IoT–fog–cloud ecosystem: A survey and taxonomy". In: *IoT* 2.1 (2021), pp. 92–118.
- [7] David Espinel Sarmiento et al. "Decentralized SDN control plane for a distributed cloud-edge infrastructure: A survey". In: *IEEE Communications Surveys & Tutorials* 23.1 (2021), pp. 256–281.
- [8] Ferdinand Brasser et al. *Trusted Container Extensions for Container-based Confidential Computing*. 2022. arXiv: 2205.05747 [cs.CR].
- [9] Annabelle Gawer, Michael A Cusumano, et al. *Platform leadership: How Intel, Microsoft, and Cisco drive industry innovation*. Vol. 5. Harvard Business School Press Boston, 2002.
- [10] Marc H Meyer and Alvin P Lehnerd. *The power of product platforms*. Simon and Schuster, 1997.
- [11] Amy J Ko et al. "An exploratory study of how developers seek, relate, and collect relevant information during software maintenance tasks". In: *IEEE Transactions on software engineering* 32.12 (2006), pp. 971–987.