



White Paper

v3.0

October 2020

www.cudos.org

Contents

1. Executive Summary	4
1.1 Directors' Responsibility Statement	6
1.2 Summary of Legal Considerations, Risks and Disclaimers	7
2. Market Overview	9
2.1 Cloud Computing Today	9
2.2 Cloud Computing Tomorrow	10
2.3 New Computing Paradigm: Blockchain	12
2.4 Why Now?	14
2.4.1 Blockchain Applications: DeFi	14
2.4.2 Physical Limits	14
2.4.3 Network Connection Improvements	15
2.4.4 Ecological Impact	16
2.5 Underutilised Resources	18
3. Cudo's Solution	21
3.1 Layer 2: CUDOS Network	21
3.1.1 How It Works	22
3.1.2 CUDOS Validator Nodes (CVNs)	23
3.1.2.1 AMD SEV Extra Security Layer	23
3.1.2.2 Delegated Staking	23
3.1.4 CUDOS Governance	24
3.1.5 Developer Journey	24
3.2 Layer 3: Cudo's Compute Platform	25
3.2.2 Supply Side	26
3.2.3 Demand Side	27
3.2.4 Developer Side	27
3.2.5 Allocation of Jobs	27
3.2.6 Pricing	27
3.2.7 Privacy	28
3.2.8 Staking	28
3.2.9 Reputation	28
3.2.10 Projected Supplier Revenue	29
3.2.11 Projected consumer savings	30
3.3 Target market	30
3.3.1 Smart Contract Blockchains	31
3.3.2 Cloud Supply Side	31
3.3.3 Cloud Demand Side	34
4. CUDOS Token	36
4.1 Token Utility	36
4.2 CUDOS Network	36

Contents

4.3 Staking in the Cudo Platform	37
4.4 Value Creation and Capture	41
4.5 Token Sale	42
4.6 Token Distribution	42
4.7 Token Schedules	43
4.8 Use of Funds	43
4.9 Security	43
5. Use Cases	44
5.1 Blockchain Compute	44
5.2 Rendering	44
5.3 AI/Machine Learning	44
5.4 Simulations	44
6. Business Development	45
6.1 Current Cudo Platform	45
6.1.1 Cudo's Ecosystem	46
7. Market Space	51
7.1 Summary	51
7.2 Layer 2/Oracle Projects	52
7.3 Other Blockchain Projects	53
7.4 Centralised Cloud Computing	54
8. Roadmap	55
9. Team	57
9.1 Core Team	57
9.2 Advisors	61
9.3 Partners	64
10. Closing Statement	66
11. Glossary	67
12. Endnotes	68

1. Executive Summary

Cudo is a scalable computing network for both blockchains and consumers of cloud services. It connects blockchain developers and buyers of cloud services (consumers) with sellers (suppliers), who can monetise their idle computing resources. Cudo's platform takes just a few clicks to install, after which the software immediately begins generating revenue. Cudo automatically and continuously generates passive income from anyone's hardware, utilising spare computing power and providing instant ROI. Blockchain developers have access to compute oracles, and enterprises can obtain computation on demand that is cheaper and more flexible than that offered by the leading cloud providers. This creates a competitive compute marketplace that offers benefits to consumers and providers.

The blockchain technology disrupted the last decade, moving from its infancy to the billion-dollar ecosystem that it is now. It is a very fast-paced environment, where new ideas and functionalities are constantly being created and implemented. Some of these innovations, such as DeFi, have started being held back by the technology's current limitations, and so they are missing the next step in terms of computing power.

Furthermore, cloud computing has been one of the defining technologies of the 21st century. By outsourcing computing and storage, users no longer have to buy and store the resources they require to conduct their operations. They can simply lease and release required resources, which has greatly increased speed, efficiency and productivity across a range of industries and verticals. Demand for cloud computing is now growing exponentially, with the dominant cloud providers spending tens of billions of dollars on infrastructure to keep pace. The new home-working policies implemented due to coronavirus are also contributing to the mass-adoption of the various cloud offerings.

The phenomenal advancement in cloud computing has not been without its challenges. The centralisation of cloud servers has led to concerns over data privacy,¹ while research has highlighted the environmental impact of hyperscale data centres. In addition,² due to physical limits in the production of microchips, the efficiencies of cloud computing compared with consumer owned devices is tailing off. The latencies incurred when transmitting data to and from the source are impractical for meeting the demands of a hyperconnected world. As George Gilder states, "Cloud computing, which was a great triumph for its time and dominated its time, is now reaching the end of the line."³

Meanwhile, consumer tech, which includes 2 billion PCs and 2.7 billion smartphones, spends the majority of the time sitting idle. The processing capabilities of these powerful devices present a huge source of potential revenue. Unlocking the value stored within these assets can provide a passive income for owners while delivering affordable on-demand compute for enterprise consumers.

Selling spare compute is the natural evolution of the sharing economy, which began with the resale of unwanted goods, and then underutilised property and automobile capacity. In the Web 3.0 era we are entering, data is the primary currency and resources are pooled, forging open cloud networks where access is certain and anyone can provide resources. Cudo will be at the vanguard of this movement, facilitating the transition to fog and grid computing, in which resources move to the edge of the cloud, creating additional demand for distributed computing and by pushing further the boundaries of blockchain's compute abilities.

Until now, decentralised computing platforms have been hampered by complexity, a narrow addressable market and poor UX. A steep learning curve has limited their use to the tech-savvy, leaving the bulk of the world's idle computing power untapped. Designed with user experience in mind, Cudo's compute platform will require no coding or complex installation and can be fully functional in minutes.

The team behind Cudo Ventures, headed by CEO Matt Hawkins, has over 16 years' experience in cloud computing having built C4L, a cloud service provider that supported 1% of the UK's internet infrastructure, with tens of millions of daily users on its infrastructure across Europe and the US. As a result, the team has consolidated an extensive network of industry contacts, including several enterprises and service providers that will become early Cudo compute consumers.

The Cudo team began operations with zero funding, covering all setup costs through the sale of C4L. Now, over three years on, Cudo Ventures has launched a professional and user-friendly platform enabling hardware owners to passively earn revenue twenty-four hours a day, seven days a week. Platform users are already earning over \$600,000 monthly. To date Cudo has built up a network of over 150,000 signed up users from over 145 countries with 20,000+ concurrent devices at any point in time, who will become early Cudo compute providers.

The final piece in the puzzle, to ensure incentive alignment and user adoption, is a well-designed token model, created with some of the industry's leading token economists, Michal Bacia, AmaZix and Outlier Ventures. As a result, Cudo compute will launch with a clear token model that will stimulate asset suppliers to become long-term network participants.

Through a token generation event, Cudo Ventures is seeking to raise \$3M to complete the Cudo compute platform, expand operations and form a distributed token-holder community, who will be encouraged to share their computing resources and be rewarded accordingly.

Thanks to its extensive knowledge of the cloud computing market and established platform, Cudo Ventures understands the needs of consumers and suppliers. In Cudo compute, it will create a network that makes renting out computing power as easy as ordering an Uber.

1.1 Directors' Responsibility Statement

The Directors of Cudos Limited have issued this White Paper and have taken all reasonable care to ensure that the facts stated in this document are true and accurate in all material respects, and that there are no other facts the omission of which would make misleading any statement in the document, whether of facts or of opinion. The Directors accept responsibility accordingly.

1.2 Summary of Legal Considerations, Risks and Disclaimers

IMPORTANT NOTICE: Please read the entirety of this “Summary of Legal Considerations, Risks and Disclaimers” section carefully. We recommend you consult a legal, financial, tax or other professional advisor(s) or expert(s) for further guidance prior to participating in the Cudos Token Generation Event outlined in this White Paper. You are strongly advised to take independent legal advice in respect of the legality in your jurisdiction of your participation in the Token Generation Event. You should note that in the Token Sale Terms and Conditions that you will be acknowledging and accepting as part of the process to participate in the Cudos Token Generation Event, you will represent that you have indeed taken independent legal advice.

Please note that this is a summary of the “CUDOS Token Terms” document which can be found [here](#) and which you must read in full before: (i) making use of this White Paper and any and all information available on Cudos Limited’s (the “**Company**” or “**Cudos**”) website at www.cudos.org (the “**Website**”) and/or (ii) participating in the Company’s token generation event outlined in this White Paper (the “**Token Generation Event**”). Any undefined capitalised terms below shall have the meaning set out in the “Legal Considerations, Risks and Disclaimer” paper. This summary should not be relied on in place of reading the “Legal Considerations, Risks and Disclaimers” paper in full. The information in this White Paper and all information available on the Website shall hereinafter be referred to as the “**Available Information**”.

The “CUDOS Token Terms” paper, the full version of which was mentioned above, applies to the Available Information. The contents of the “CUDOS Token Terms” paper outlines the terms and conditions applicable to you in connection with (i) your use of any and all Available Information; and/or (ii) your participation in the Token Generation Event, in each case in addition to any other terms and conditions that we may publish from time to time relating to the Token Generation Event (such terms hereinafter referred to as the “**Terms**”).

This White Paper states the current views of the Company concerning the Cudos Compute Platform and related matters. The Company may from time to time revise this White Paper in any respect without notice. The information entered in this White Paper is indicative only and is not legally binding on the Company or any other party. This document is for informational purposes only and does not constitute and is not intended to be an offer to sell, a solicitation of an offer to buy, or a recommendation of: (i) the Company, (ii) an investment in the Cudos Compute Platform or any project or property of the Company, or (iii) shares or other securities in the Company or any affiliated or associated company in any jurisdiction.

The information set forth in the “CUDOS Token Terms” paper may not be exhaustive and does not imply any elements of a contractual relationship. While we make every reasonable effort to ensure that all Available Information is accurate and up to date, such material in no way constitutes professional advice. Individuals intending to participate in the Token Generation Event should seek independent professional advice prior to acting on any of the Available Information.

The Company does not recommend purchasing Tokens for speculative investment purposes. Tokens do not entitle you to any equity, governance, voting or similar right or entitlement in the Company or in any of its affiliated companies. Tokens are sold as digital assets, similar to downloadable software, digital music and the like. The Company does not recommend that you purchase Tokens unless you have prior experience with cryptographic tokens, blockchain-based software and distributed ledger technology and unless you have taken independent professional advice.

Citizens, nationals, residents (tax or otherwise), green card holders and/or Restricted Persons of any Restricted Jurisdiction shall not process the Available Information and are prohibited from participating in the Token Generation Event or the purchase of Tokens or any such similar activity.

In no event shall the Company or any current or former Company Representatives be liable for the Excluded Liability Matters.

The Company does not make or purport to make, and hereby disclaims, any representation, warranty or undertaking in any form whatsoever to any entity or person, including any representation, warranty or undertaking in relation to the truth, accuracy and completeness of any of the information set out in the Available Information.

You should carefully consider and evaluate each of the risk factors and all other information contained in the Terms before deciding to participate in the Token Generation Event.

This White Paper may be translated to different languages but in the event of a conflict between documents, the English version of the White Paper will prevail.

2. Market Overview

2.1 Cloud Computing Today

Cloud providers in possession of hardware such as servers and data centres rent their computing capacity to cloud consumers. Although thought of as a recent phenomenon, cloud computing can be traced back to the 1960s, when the concept of an “intergalactic computer network”⁴ was proposed. The first manifestation of this was ARPAnet, a forerunner to the internet, which offered a number of advantages over the bulky mainframe computers that emerged in the 1950s.

As technological advances shrunk computers, which were initially the size of a small house, a distributed network began to form, with the initial goal of facilitating scientific research and military communication. A client-server computing model took shape and persisted until the 1990s, when the internet as we know it emerged. The number of networked devices proliferated, aided by reduced hardware costs, and dramatically increased processing power and data speeds. In the late 90s, Salesforce launched the first Software as a Service (SaaS) applications and the modern cloud computing era was born, though it would be a number of years until “the cloud” entered common usage.

The dawn of the 21st century saw consumer demand for ever more portable yet powerful devices whose design mandated offloading computation to the cloud. Mobile computing exploded, enterprise software rapidly migrated to the cloud, and data centres multiplied and expanded. Software companies such as Amazon, Google and Microsoft introduced cloud computing to end users with a suite of powerful applications, and dedicated cloud verticals began to form including Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). The latter is now a \$39B industry,⁵ while the global public cloud service market is set to grow from \$182B in 2018 to \$331B by 2022.⁶

In 2018, the number of large data centres operated by hyperscale providers – primarily Amazon, Microsoft, Google and IBM – rose by 11%.⁷ In 2017, the top 24 providers spent \$75B on data centres and server farms.⁸ Google invested \$13B in cloud computing through 2019 including thousands of new staff, nine new offices and six data centres across 14 states.⁹ 2020 and the COVID-19 pandemic have accelerated the mass-adoption of cloud computing, which has arguably become a basic need of many businesses.¹⁰ Home working started as a necessity during the pandemic, but many sectors are likely to maintain it afterwards as a viable alternative.¹¹

The history of consumer level computing can be seen as a race to create increasingly compact and powerful devices. This macro narrative has played out for the past 70 years as the industry has oscillated between periods of computational centralisation and decentralisation. The current hub and spoke system, in which centralised data centres handle much of the web’s storage and processing, has been in place for almost two decades. History would suggest that the pendulum is due to swing once more, to a distributed model that meets the demands of next generation networked devices.

Today, vast amounts of spare capacity remain underutilised across a range of consumer hardware while expensive, and energy intensive data centres serve devices which sit idle for a substantial amount of the time. This situation is extremely inefficient. The combined computing potential of these underutilised consumer devices is equivalent to hundreds of large-scale data centres and could, if properly mobilised, fulfil a large portion of the world’s computing needs. Due to the current lack of ability to fill this spare capacity, much of it remains idle.

2.2 Cloud Computing Tomorrow

The maturation of cloud computing, coupled with the evolution of networked devices has given rise to a number of new cloud-based models. Although niche at present, just as the centralised cloud was 20 years ago, there is compelling evidence that these new cloud-based systems will play a defining role in the future of computing. Edge computing functions by processing data in close proximity to its origin via micro data centres rather than sending it back and forth between the data source and the cloud data centre, to increase efficiency.

Fog computing is essentially the midway point between cloud and edge computing, with data processed at multiple points within the network. The fog computing market has been predicted to grow fourfold by 2022, reaching \$18B.¹² Finally, grid computing involves processing power, memory and data storage being shared with every other computer on the network, creating a distributed system of non-interactive workloads.

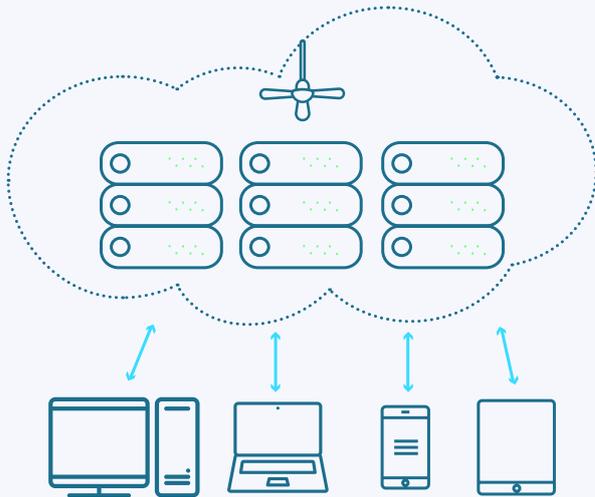
In a grid system, devices in a network send tasks that require a large number of processing cycles and involve large amounts of data to a central entity. This entity breaks the tasks up into fragments and distributes them to a range of connected computing hardware. Each of these devices solves their allocated fragment and returns the completed task to the central entity, which re-constitutes the fragments into a full, completed job and returns it to the origin device. This system is secure since no device in the network save for the central entity has access to the complete task or result. It is also fast since multiple devices work in parallel to complete segments of the job, and efficient because it can make use of spare capacity from otherwise idle devices attached to the network.

Since the data collected from IoT devices is often unstructured, machine learning techniques in edge systems are required to make real-time decisions, and thus low latency is imperative. The growing power of IoT sensors, mirroring the growth of mobile processing power, will produce massive amounts of data, increasing the need for edge computing. Edge computing will enable applications such as vehicles in geographical proximity communicating with one another in real-time, effectively creating distributed P2P networks that can rapidly process and transmit data.

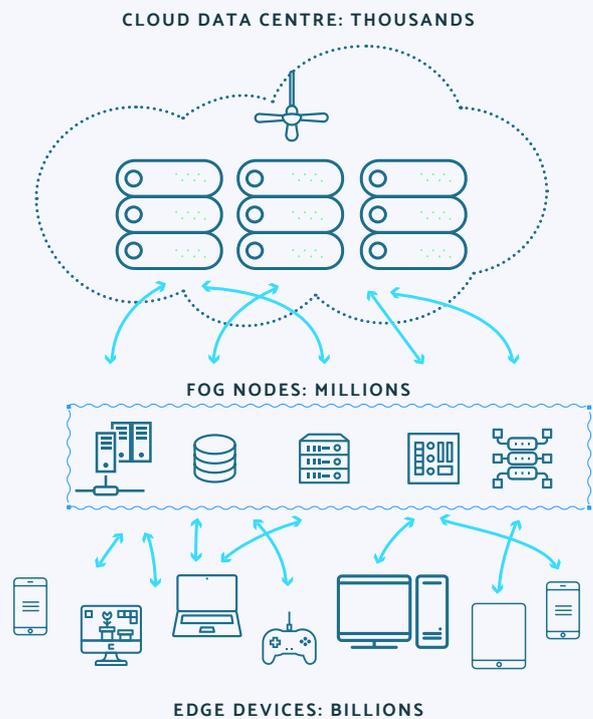
The rise of drones, self-driving cars, robots and IoT devices that are reliant on low latency processing will drive more compute demand to the edge. By 2025 there will be more than 75B connected IoT devices that need to send and receive computations in real-time.¹³ As billions of devices join the edge of the network, computing power will be forced to follow, giving rise to an era of distributed computing. This shift has been described by a16z's Peter Levine as "one of the biggest transformations to occur" to the computing landscape.¹⁴

While grid and fog computing create an interconnected network of devices, these solutions share one key thing in common with the current cloud computing: trust in a single entity controlling them all or, in other words, centralisation. In 2008, a new technology called blockchain emerged, launching a currency that did not need to trust any single entity: Bitcoin.¹⁵

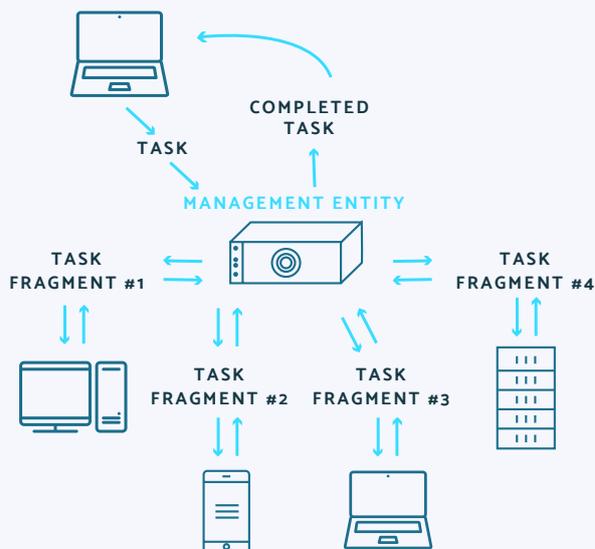
Cloud



Fog



Grid



2.3 New Computing Paradigm: Blockchain

In 2008, Satoshi Nakamoto, whose identity is yet to be revealed,¹⁶ proposed the first implementation of a decentralised, trustless digital currency that solved the double-spending problem: Bitcoin. Bitcoin uses what is known as blockchain, a distributed and decentralised public ledger.

Simply put, Bitcoin's blockchain is formed by nodes, and each node has a copy of the ledger. New transactions in the network are grouped into blocks, and transaction and block uniqueness is guaranteed by cryptography algorithms. To be more precise, Bitcoin (and most other cryptocurrencies out there) use what is known as Proof of Work (PoW): each new block has a unique hash connecting it to the chain, and all nodes are constantly trying to calculate that hash. The winner gets rewarded in Bitcoin, broadcasts the block to the network and starts to find the hash for the next block (what is known as mining).¹⁷

Controlling most of the newly mined blocks would imply controlling the network, and being able to alter data. This is known as a 51% attack, and may pose a problem to small cryptocurrencies. However, for bigger networks, as has been the case for Bitcoin and many others for years, this attack is unfeasible, as the computing power required to successfully deliver such an attack would need to be massive.¹⁸

While Bitcoin's network is used to trade the native cryptocurrency, developers around the world soon started to build on top of the blockchain technology, to expand its functionality range. Most noticeably, Vitalik Buterin proposed what is today the second biggest cryptocurrency: Ethereum. Ethereum implements smart contracts, which run on the Ethereum Virtual Machine. The biggest revolution about these smart contracts is that they are written in a Turing-complete language.

Turing completeness means that Ethereum's smart contracts can run any kind of algorithm, unlike Bitcoin which is mainly used to transact cryptocurrency. This opened a new landscape for blockchains, as it gave developers the possibility to run any kind of workload on-chain. This new functionality is limited though: Ethereum's internal transactions are priced using gas,¹⁹ which limits the amount of transactions and compute that can be run on-chain. Also, PoW greatly limits the number of transactions that any blockchain can do per second, and is a very inefficient system. For instance, in July 2019 Bitcoin was reported to be consuming more electricity than Switzerland.²⁰

Another expression of this gas issue can be seen very clearly when looking at the computing, bandwidth and storage costs in these platforms.²¹ As of October 2020, 1GB of storage in Ethereum is counted in tens of thousands of millions of dollars, and in EOS in hundreds of thousands of dollars. On the other hand, the same storage size in Amazon Web Services (AWS) only costs cents, and the amount of transactions per second it supports is many orders of magnitude higher.

Although, or maybe thanks to these scalability issues, Bitcoin and blockchain revolutionised the developers and researchers ecosystem, by introducing a very interesting area of research²² to work and develop on. However, blockchain quickly became much more than that. One Bitcoin's market price was around \$400 by the start of 2016, it was almost \$1k by the start of 2017 but then skyrocketed to almost \$20k by the end of that year. Even if it quickly went down after that, the price of one Bitcoin has rarely gone below \$6k at any point after that, and has been oscillating around \$10k in the last months.²³

Therefore, Bitcoin, Ethereum and all the other ecosystems of cryptocurrencies (commonly called altcoins) have also become investment assets, welcoming a much wider audience into the cryptocurrency world. Also, thanks to the possibility to easily launch tokens within smart contract platforms such as Ethereum, many development projects got the opportunity to research and move the field forward from within, receiving support through their currencies and tokens within the crypto community.

While some projects made it to mainstream media but did not bring much value even though they flagged underlying issues,²⁴ others such as crypto.com, EOS or Tezos have been introducing cryptocurrencies to an ever wider audience and providing services in many different areas. Furthermore, the whole ecosystem has been evolving very rapidly in many directions. Most recently, decentralised finance (DeFi)²⁵ has become a very hot topic, growing from just under \$700M in January 2020 to over \$11B at the end of September.²⁶ However, the aforementioned issues still dampen its functionality.

In order to overcome the issues introduced by PoW, many projects have started to look for alternatives. Proof of Stake (PoS) is one of the main candidates, which has already been implemented by projects such as Cosmos or Polkadot, and is planned in Ethereum for the 2.0 release, by using sharding. The main idea behind PoS is that nodes validate blocks based on the amount of tokens they have staked, rather than competing with each other to win the PoW race.

In a different direction, projects like Algorand have an improved Byzantine Agreement system, which greatly improves the transaction throughput but does so at the cost of not having Turing-complete smart contracts. What many of these new projects have in common (or are missing) is what is known as a layer 2: a separate network of nodes which can run code and validate transactions off-chain.²⁷

Similarly, recent projects have tried to connect blockchains to the outside world, via oracles. Examples include Chainlink, a very interesting project that provides data feeds for currency values. The key innovation that Chainlink is providing is that these data feeds are as decentralised as the requestor wishes: the smart contract requesting the job can choose how many Chainlink nodes to use in order to receive the currency value. However, gas costs are still very high, and so there is no definitive solution to the scalability problem.

2.4 Why Now?

On the blockchain side, there is a clear need for a scalability solution that can be implemented across different blockchains, and that allows blockchain developers to continue pushing the field forward. On the centralised cloud side, issues on both the supply and demand side threaten to stifle the growth of the cloud computing industry, driving businesses to more efficient and flexible solutions.

On the demand side, enterprises waste 35% of their cloud spend on average through ineffective management and optimisation.²⁸ This translates to higher expenditure than necessary with providers like AWS and has been calculated to cost companies \$62B annually.²⁹

Such is the gravity of the challenges affecting the industry that numerous experts have predicted that a radical shake-up of the cloud business model is inevitable. Given the recent and lightning-fast growth of the blockchain industry a connection between both worlds is not only necessary, but also inevitable.

2.4.1 Blockchain Applications: DeFi

While the biggest smart contracts blockchain, Ethereum, is moving towards Ethereum 2.0 to solve the scalability issues, it is doing so slowly, and the crypto community is moving faster than it. This is very clear looking at Ethereum's transaction costs: since DeFi has risen in the market, transaction costs have increased by an order of magnitude.³⁰

Furthermore, DeFi applications are limited in scope, as is the rest of the ecosystem until viable solutions are found. The technology is currently at a turning point, after which it will evolve even faster towards mass adoption. There are many projects out there offering funds and grants for blockchain research,³¹ and the whole crypto ecosystem has been steadily growing for years to reach this point. Therefore, it is the perfect point in time for professional cloud enterprises to start collaborating and working within the blockchain world.

2.4.2 Physical Limits

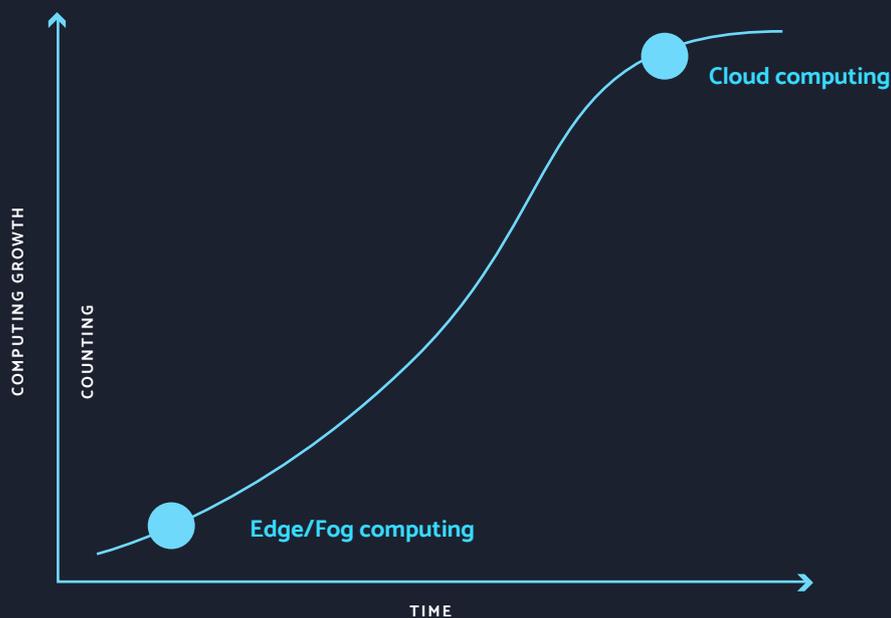
On a more general computing aspect, many of the challenges the industry is facing are not unique to the cloud, but rather affect the entire tech sector. Dennard scaling – the principle that as transistors become smaller they use less power, thus reducing the heat generated and allowing them to be packed closer together – has reached a breaking point. This is on account of the small leakage current that heat transistors possess.

As the number of transistors for a given area increases, the leakage current dominates the gains due to Dennard scaling. In other words, the power consumption to the silicon cannot scale in line with scaling of the processors. This results in what is known as “dark silicon”³² where the cores in a multicore processor cannot all function simultaneously due to power constraints.³³

Although the production size of transistors is decreasing, this no longer provides the equivalent increase in computing power that it used to because only part of it can be fully operational at any point in time to prevent overheating.

As a result, improvements in CPU clock speeds have been decelerating since 2005³⁴ and data centre efficiency advantages have begun to tail off relative to other forms of computing infrastructure.

Computing adoption follows a conventional S curve, with cloud computing now approaching the asymptote of that curve and slowing while other verticals are still rapidly accelerating. The specialist hardware operated by cloud providers has reached a cap where additional capacity can only be attained through the construction of new and larger data centres, which introduce their own problems.



2.4.3 Network Connection Improvements

Until recently, poor network connections differentiated homes and businesses from data centres. Commercial and domestic connections have historically lagged far behind in terms of reliability and speed. This has prohibited domestic devices from providing compute due to intolerably slow latency and the risk of failure due to dropped connections. However, the connection gap is quickly closing. As of 2018, 94% of UK homes and businesses have “superfast” broadband,³⁵ 6% have access to full-fibre connections,³⁶ and 91% of the UK’s landmass has access to good 4G mobile coverage from at least one operator.³⁷ The 47 least developed countries in the world (LDCs) now all have 3G internet with 60% of their population covered by a 3G network on average. The UN projects that these LDCs will achieve on average 97% mobile broadband coverage by 2020.³⁸

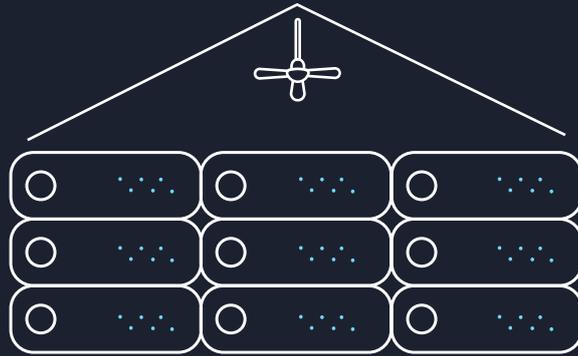
These rapid and widespread improvements in connectivity have profound implications for the future of computation. Where connections were previously unreliable and download speeds were slow, advancements in this field have led to a large-scale network of users with fast and reliable connections. Now, with rapidly improving network connections, it has become feasible for these entities to supply compute to other users.

2.4.4 Ecological Impact

The carbon emitted by data centres is enormous, with cloud operations consuming gigawatts of energy per year. The world's 430 hyperscale data centres are set to be supplemented by a further 132 that are under development, causing a commensurate increase in energy consumption to keep up with consumer demand.³⁹ Although hyperscale cloud data centres typically operate more efficiently than their independently operated counterparts,⁴⁰ the cloud model still incurs a high environmental cost.

Up to 50%⁴¹ of the energy required for cloud computing is used for air conditioning and other non-computing power demands.⁴² Though some of the energy used by cloud providers is sourced from renewables,⁴³ a great deal is derived from non-renewable sources.^{44, 45} Google's carbon footprint was estimated to be over 1.6M tonnes CO₂ equivalent in 2013, the bulk of which was from data centres.⁴⁶ For context, a typical passenger vehicle emits around 4.6 tonnes of carbon dioxide per year.⁴⁷ It would take a car over 380,000 years of use to meet this level of CO₂ output.

All forms of computing are bound to consume large amounts of energy at scale, but centralised cloud computing is particularly inefficient due to its inflexible architecture and remoteness. Furthermore, the ecological impact of centralised cloud computing extends beyond mere energy consumption.⁴⁸ Much of the deleterious environmental effects are incurred at the point of manufacture. The manufacturing process is highly carbon intensive. From the initial extraction, production and transport of raw materials, the creation of components and subassemblies to the final assembly, the process of constructing a computing device involves constant energy consumption. Beyond hardware creation and variable costs such as electricity, the construction of data centres involves vast carbon consumption. This in turn is dwarfed by the enormous carbon footprint accrued by data centres once operational.



1MW DATA CENTER



Electricity
177M KW-hr



Water
227M L



Plastic
15K Kg



Copper
66K Kg



Aluminium
33K Kg



Solder
5K Kg



Steel
170K Kg



Lead
10K Kg



Concrete



Masonry



Metals



Wood



Chemicals



Glass



Tiles



Water



Plastic



Insulation



Tar

It is difficult to calculate the exact CO2 cost of sourcing and constructing all of these components, but an estimate can be made. The materials required to construct a 5,700 sq. ft. office facility have been calculated to have a carbon emission contribution equivalent to 128.3 tonnes of CO2.⁴⁹ AFCOM's 2019 State of the Data Center Industry report found the average size of all data centre space to be between 180,000 and 240,000 sq. ft.⁵⁰ Assuming a 210,000 sq. ft. average, a data centre is around 38 times larger than the office space and therefore has a carbon emission contribution equivalent to roughly 4,875.4 tonnes of CO2. This number is purely based on the embedded CO2 in the building materials required to construct the building shell and does not take into account the carbon costs of sourcing, extracting and transporting these materials or the construction process. It would take approximately 1,060 years of car journeys to reach this carbon cost.

Contained within each data centre is an array of operations hardware including Computer Room Air Conditioning (CRAC) units, heat rejection devices such as dry coolers and air cooled chillers, equipment to pump chilled water between CRAC units and dry coolers, Uninterruptible Power Supply (UPS) modules, generators, storage units, electrical primary switch gear, servers, Storage Area Network (SAN) hardware, Wide Area Network (WAN) hardware, Local Area Network (LAN) hardware, Power Distribution Units (PDUs), and racks to hold all of this hardware as well as miles of cabling.

2.5 Underutilised Resources

There are approximately 2 billion PCs,⁵¹ 2.7 billion active smartphones,⁵² 138 million Playstation 4 and Xbox One gaming consoles,⁵³ and 30 billion IoT devices that are capable of performing some degree of data processing.⁵⁴ Estimates place the amount of data centre IT hardware that is unutilised at upwards of 50%.^{55, 56} Workplace devices such as office computers are utilised even less. In the UK, office computers are used for an average of 6.5 hours per day and unutilised close to 75% of the time.⁵⁷ Consumer hardware like laptops and desktop PCs is, on average, utilised even less. One study found that half of all computers are operational for less than five hours per day.⁵⁸ Such devices are underutilised 80% of the time, generating zero value for the asset owner.

The same is true of smartphones, which are used intermittently throughout the day, but whose total percentage of utilisation is minimal. While the owner is asleep, these devices are fully charged and sit idle with their computing potential unrealised.

A major inefficiency in enterprise computing is capacity hoarding. KPMG explains: "Capacity hoarding occurs primarily as a result of organizations not being able to accurately forecast their future capacity needs coupled with the typically long lead times to provision new infrastructure. Rather than risk the potential of not having sufficient capacity when needed, excess capacity is held back."⁵⁹ This surplus computing capacity is rarely used fully and constitutes another segment of computing assets locked away and underutilised.

Even cloud platforms underutilise a significant portion of their compute capacity. Current estimates show 35% of cloud capacity is being wasted, partly due to the requirement to retain spare capacity for on demand workloads. This surplus compute cannot be resold to contracted customers, and needs to be kept for the peaks and troughs of demand. This translates to higher-than-necessary costs to the tune of \$62 billion annually.⁶⁰

Building a data centre is incredibly costly, energy intensive, and creates a significant carbon footprint. If the completed data centre is not filled with the maximum amount of computing hardware to perform the maximum amount of computation possible then it is, by definition, being underutilised so the ratio of construction costs-to-profits becomes higher.

The amount of compute hardware in a data centre must match the maximum amount that can be supported by the secondary systems used, such as Uninterruptible Power Supply (UPS) systems and cooling systems, to be optimally efficient. Where the amount of compute hardware used falls below the maximum supported by the secondary systems, all of these secondary systems become less efficient since they are not being utilised to their full potential.

The spare capacity that is locked in home computers, smartphones, servers and game consoles is untapped at present due to technical challenges, logistics and lack of economic incentives. Since this hardware has already been purchased, the only additional expenditure required for it to operate at capacity is variable costs, namely the cost of electricity.

Centralised cloud computing, in comparison, incurs ongoing capital expenditure such as costs of hardware renewal, cooling systems, the facilities used to house the hardware, staff and administrative costs. As such, spare capacity in personal devices can be deployed much more efficiently and cost effectively than cloud computing on an ad hoc basis.

Mobilising existing hardware and exploiting the full potential of domestic computers would also serve to reduce carbon emission. Research indicates that distributed computing consumes 14%-25% less energy than fully centralised systems, partially through not being reliant on intra-data centre networks and industrial cooling systems.^{61, 62}

While it is true that network suppliers must pay the cost of their energy consumption, this will be more than compensated by the rewards earned for providing computation. Profits will naturally vary depending on the hardware used and the variable costs of running it (e.g. electricity and maintenance). However, users can be expected to profit in every scenario, as detailed in section 3.2.10.

The utilisation of spare capacity has been popularised through the efforts of companies such as Airbnb, Uber and Lyft. Businesses and end-users are comfortable with the notion of repurposing idle capacity to yield greater efficiencies, cost savings and monetisation opportunities. Computing is the next industry to be disrupted in this manner through the introduction of a business model whose time has come.



Extraction of raw materials



Transport of raw materials



Manufacture of components



Subassemblies



Product manufacture



Final assembling

227 - 270 Kg CO₂



Standard laptop used for 8 hours per day

3. Cudo's Solution

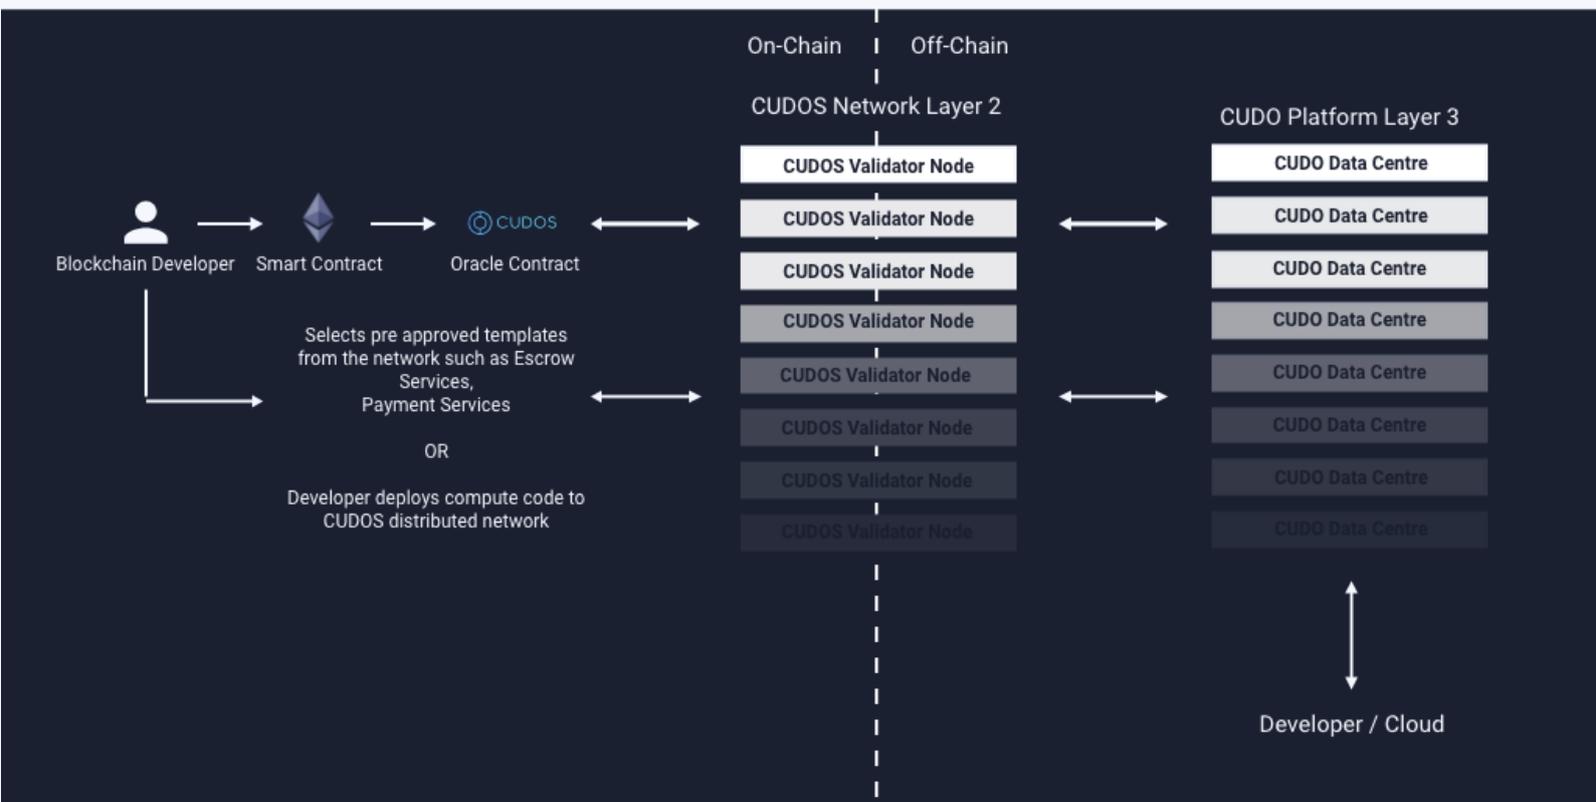
Cudo provides scalable compute solutions, and has identified two very clear opportunities: to provide secure layer 2 computing capabilities to blockchains, and to improve the offerings on cloud services. However, joining the dots Cudo has envisioned a step further: connecting the two, providing for the first time a novel layer 3 to blockchain networks, in order to allow them to run any kind of compute jobs, by seamlessly integrating its cloud solution to the blockchain one.

3.1 Layer 2: CUDOS Network

The CUDOS network provides a layer 2 to various blockchains, in order to solve the current scalability issues. The CUDOS network comprises a set of smart contracts that act as compute oracles, connecting blockchains to secure external sources in a decentralised way. By allowing smart contracts to request off-chain computation, CUDOS offers a Turing-complete solution which can run code in any language, such as Python, Go or C++.

3.1.1 How It Works

The connection to the outside world is done via the CUDOS Validator Nodes (CVNs), which are one of the main pillars of the CUDOS network. Blockchain developers are able to choose through the CUDOS smart contracts the level of decentralisation that they want, by selecting how many CVNs should receive and run the desired code. This gives the power and freedom to developers to choose and select the most appropriate settings and configuration for each job.



Once each CVN has run the job, they all share the result and run the validation and consensus piece. Consensus in the CUDOS network can run either on-chain or off-chain, to further reduce the network's transaction costs, at the will of the developer requesting the work. This system gives the control back to developers, not only allowing them to choose the level of decentralisation they want, but also controlling the cost for the overall request.

3.1.2 CUDOS Validator Nodes (CVNs)

CVNs are the backbone of the CUDOS network. Initially, the CUDOS network will support a maximum of 100 service providers, with a maximum number of 1,000 CVNs. Each node needs to stake 2,000,000 CUDOS to become a CVN. The first batch of validators has hardware boxes validated by Cudo, in order to ensure that the specs allow all the intended work to run smoothly, and get rewarded accordingly. These work and rewards include:

- 24x7 revenue provided from the layer 3 cloud side (see section 3.2)
- Revenue from completed blockchain compute jobs through the CUDOS network
- Staking rewards for the contribution to the network
- Discounts on the fees paid

Another benefit from the CUDOS network is that, on top of the security added by its decentralised nature, all CVN's will have hardware with encryption capabilities, adding an extra layer of security to the workloads for all sides.

3.1.2.1 AMD SEV Extra Security Layer

AMD offers two main security features: Secure Memory Encryption (SME) and Secure Encrypted Virtualisation (SEV). SME comprises a hardware memory controller, which encrypts and decrypts data as it travels between the DRAM and the CPU.

SEV protects guest VMs from a malicious hypervisor, by running in a secure enclave where data cannot be accessed by any other software in the machine. In order to avoid the hypervisor accessing the register contents of a VM right after it stops, SEV Encrypted State (SEV-ES) encrypts all CPU register contents when a VM stops running.

SEV encrypts Virtual Machines (VMs), in order to protect the data running on them from attacks such as physically removing a non-volatile memory module. It also protects the data from other VMs running on the same machine, or from the hypervisor hosting it. To protect the guest register state from the hypervisor, SEV-ES enabled VMs have control over the exits to the hypervisor.

Encryption keys are tagged to each secure VM, and that VM's data is restricted to use that tag. These encryption keys provide proof that SEV is supported on the hardware and that a guest VM has been deployed securely with SEV. Furthermore, SEV uses a remote attestation system. A set of keys are created by the SEV firmware in order to establish a secure channel between the guest owner and the SEV platform, and to prove that the guest VM is in the desired state.

There is a special system-on-chip which creates the ephemeral encryption keys, and an AMD server which verifies the authenticity of the identity key of each SEV firmware. As CVNs have AMD EPYC chips which support SEV, once the blockchain workloads start running on these chips, this technology adds this extra layer of protection to both transacting parties.

3.1.2.2 Delegated Staking

Users willing to contribute to the CUDOS network but not desiring to run a full CVN can delegate their stake to a CVN. By doing that, they will be rewarded for their contribution to the network, similar to how the CVNs which they support are.

By receiving more delegated staking, the validator nodes will increase their trust score. The trust score is a metric to guide blockchain developers when choosing which nodes to use, by giving an indication of different characteristics of each node, including:

- Staking score
- Reliability score
- Availability score
- Security score

3.1.4 CUDOS Governance

CVNs will have control over the governance of the CUDOS network. The trust scores will be used to calculate a weighted average score, which will determine the weight of each node in the decision making process. By limiting the total amount of CUDOS tokens staked and the delegated staking that each CVN can receive, as well as introducing weighted probability algorithms to distribute the selection of nodes when developers do not explicitly choose them, the network is expected to have an even distribution of stakes among a sufficiently large number of nodes.

Each CVN will be free to choose how to involve users delegating their stake to them. For example, they could choose to regularly propose polls about the most discussed topics within its delegators, to then propose changes on the network based on the outcomes. In order to ensure a smooth start of the platform and to prevent abuse, Cudo will control the governance of the network in the initial alpha stages. When validators have earned a high enough score and the platform has reached enough maturity and usage, governance will be fully handled by the nodes.

3.1.5 Developer Journey

From a blockchain's developer perspective, the CUDOS network is a set of smart contracts which can be called from their desired contract. Developers are free to choose the level of decentralisation or, in other words, to how many and which validators nodes they want to send the job request. In order to avoid on-chain transaction costs, the exchange of actual data is done off-chain.

Similarly, developers are free to choose the consensus method, which can be run on-chain, selecting one of the predetermined CUDOS methods, or off-chain with custom code, especially for more complicated consensus decisions. To run code and the consensus part, developers can either write their own, or choose an app from the CUDOS marketplace, which will offer already written and validated code for some of the most common use cases.

3.2 Layer 3: Cudo's Compute Platform

Cudo Ventures has created a suite of user-friendly tools that automate the process of reselling unused computing power. Its flagship product for the computing market is a collaborative consumption network that allows users to sell spare computing power to consumers. Cudo will unlock the potential in billions of devices to create a highly efficient, secure and cost-effective computation market that generates value for all participants.

Unlike existing distributed compute projects, Cudo's products are effortless to set up and use. After installing the software on their device, asset owners can begin earning in just a few clicks. Once a user has signed up, Cudo's software automatically runs local tests to assess the capabilities of the connected device. Cudo has also crowd sourced data for each hardware type across thousands of different types of hardware, which enables the platform to automatically know the best workloads for the user's hardware. Using a combination of local tests and this crowd sourced data, Cudo software selects the compute jobs that will be most profitable for each piece of connected hardware. Jobs are then allocated and operations begin. The entire process occurs automatically, requiring no significant input from the user.

Cudo's smart management layer monitors a variety of network metrics including cryptocurrency market prices, the mining difficulty of a range of coins, the size and type of computation jobs submitted to the network, and the types of hardware connected to the network and their performance metrics to identify the most profitable operations for each connected device. Based on this analysis, Cudo's software automatically adjusts and optimises settings and selects appropriate jobs. This ensures that each device is always being used as effectively as possible to provide fast and efficient job completion on the demand side and maximum profitability for compute suppliers.

Where there are no compute tasks available or where cryptocurrency mining becomes more profitable than completing the compute tasks available, Cudo will automatically allocate mining operations to machines that would otherwise sit idle. With automated switching from compute to mining, hardware owners are guaranteed to constantly generate revenue.

Through focusing on user-side simplicity, Cudo aims to make computation universally accessible, overcoming one of the biggest drawbacks to distributed computing and allowing anyone to join the network and trade computation.

Cudo consumers will include a range of enterprises, NGOs, developers, service providers and institutions that require significant computational resources on a regular or sporadic basis. The primary benefit to these entities will be the significant cost savings compared to traditional cloud computing with the additional benefits of increased compute coverage, lower latency to compute workloads and increased on-demand scalability of compute. Cudo has already signed MoUs with several large-scale demand side enterprises and is currently in talks with a number of major game console and game software & hardware providers.

Cudo compute suppliers include any entity in possession of underutilised computing assets. This encompasses a broad range of hardware including but not limited to cryptocurrency miners, servers, gaming PCs, standard PCs and laptop computers, games consoles, IoT devices and mobile devices. Cudo compute's compatibility with such a wide range of devices means that anyone can participate and become a supplier.

The benefits to asset suppliers, meanwhile, will be in the form of new or increased revenue. Cryptocurrency miners are among the first supply-side adopters of Cudo compute due to the powerful computational assets they possess. Cudo already has over 20,000 concurrent devices at any point in time and over 30,000 active devices each month that are expected to adopt Cudo compute jobs. Their high powered CPU and GPU devices are ideally suited to perform distributed computing.

Calculations show that at certain times, depending upon various factors such as cryptocurrency hash rates and prices, miners can increase their profits by 2-4x by switching to distributed computing. One of the most common concerns among crypto miners is the uncertainty of generating a sustainable income. Distributed computing will offer miners an alternative and more stable source of income, ensuring that they are no longer reliant on the health of the cryptocurrency markets for revenue generation.

Consumer Benefits

• SMART MANAGEMENT

Unlike many cloud services, Cudo allows consumers to purchase as much or as little computing power as they need whenever they need it without paying the high premiums for on-demand. This eliminates the need to estimate capacity and risk wastage. Cudo provides flexible computation on demand.

• GREATER EFFICIENCY

A significant proportion of enterprise IT expenditure costs goes on human capital including maintenance and administration.⁶³ Cudo eliminates these costs altogether. By utilising Cudo, enterprises can optimise workflows, improve decision making, and reduce problems arising from human errors.

• REDUCED COSTS

Consumers are calculated to save over 75% on average by switching from their current cloud provider to Cudo. The first consumers who are lined up to use the service will be saving 70% in CPU video rendering, 75% in university research, 85% in GPU video rendering and up to 95% in simulations.

3.2.2 Supply Side

The Cudo platform requires one click to install and is free to download. The initial version of the software has been developed for Windows, Linux and Mac. Since the software is automated for maximum efficiency and profitability, the user will not need to perform any actions beyond installing the app, adjusting a few simple parameters and ensuring that the hardware is connected to the network. Thereafter, suppliers can simply 'set and forget,' leaving Cudo to run in the background and generate revenue.

Suppliers can set their own asking price for their computational power anywhere above the baseline cost. They can also set their cost of electricity and % above electricity, minimum revenue or % above mining.

Upon installation, Cudo runs a scan to determine the hardware specification before proceeding to allocate suitable jobs for the device to complete. Machine learning techniques are used to automatically evaluate the types of jobs and workloads that are best suited to each device connected to the network. The software performs continuous benchmarking on the hardware to ensure that the most appropriate jobs are allocated to each device. In doing so, Cudo ensures that the network is always running at optimal efficiency with each job completed as quickly as possible via appropriate hardware. This results in the timely return of completed jobs for those requesting the work and a quick job turnover rate for suppliers who can profit greatly from the rewards earned. Automated benchmarking on network connectivity is also performed to ensure the job can be completed on time without disruptions. Following Cudos' decentralisation vision, the aim is to decentralise part of this process by completely merging it into the CUDOS network.

3.2.3 Demand Side

Cudo acts as a free market system where the computational price offers set by suppliers are matched to ensure that the consumer gets the best price available at any given time in light of the requirements of the task submitted. When a consumer requires computational power, they submit their task or the resources required to the Cudo network. Cudo offers suppliers recommended pricing values to assist with selling their services at a fair market price. Cudo's intention is to decentralise this recommended pricing as well by having smart contracts manage a mature billing method. Cudo has already signed MoUs with a group of key demand side companies and is in talks with a number of others including several major game console providers.

3.2.4 Developer Side

Developers can build their own applications to be listed in the Cudo marketplace app store. For example, SaaS providers can build applications, and in the future dApps, that are designed to facilitate specific types of computation. Game developers can create applications that make use of game consoles, harnessing their processing power when these devices would otherwise sit idle. Initially, the marketplace will be centralised but once Cudo starts their path to decentralisation, it will be curated by the community.

3.2.5 Allocation of Jobs

Jobs will be allocated to different networked machines according to a number of parameters. These include:

- Task priority
- Price range accepted for the job
- Requirements of the job sent
- Supplier reputation
- CUDOS availability of the supplier (staking)
- Location
- Security level

For a detailed breakdown of jobs allocation within the Cudo network, see the Cudo Technical Paper.

3.2.6 Pricing

Compute prices will be quoted in fiat currency for the purposes of pricing stability. Payment can be made in fiat or in a range of crypto assets. Cudo will charge a fee of 30% for compute jobs completed on the network, the remainder being paid to the asset suppliers for their services.

Pricing will work like an auction system, so that during peak demand, suppliers will earn more and consumers will need to pay more to have the task completed within a particular time frame.

3.2.7 Privacy

Cudo will handle various types of workloads with various privacy requirements. Some data may require advanced privacy measures. Cudo runs in different environments, some of which are inherently secure. These include ISO 27001, SAS 70 and PCI-DSS security approved data centre environments. Consumers with high-security requirements can select to only run their workloads in those environments. This will incur higher costs but will ensure privacy and data integrity. Users with workloads that do not require advanced privacy (e.g. public data workloads) can choose lower cost environments such as end user devices. Users will also have the option to run fully encrypted workloads. This is where the owner of the data holds the encryption key.

Secure hardware including AMD SEV and Intel SGX chips will be utilised, with chip encryption delivered through enclaves, creating a secure environment that requires a token to access. Encryption will also be used for distributed storage via the likes of IPFS.

3.2.8 Staking

For a supplier to be able to accept a job, they must stake a certain amount of CUDOS tokens. The amount of CUDOS required will be proportional to each supplier's contribution to the network, with lower minimums for suppliers with accrued reputation. Thus, only suppliers who stake large amounts of CUDOS will be eligible to accept large jobs. Since suppliers who stake significant amounts accept a significant risk, consumers can use this metric to make assumptions about which suppliers are likely to be trustworthy.

Suppliers do not need to make any initial investment to begin staking. Upon downloading the Cudo platform, Cudo will start running and can optionally earn users enough CUDOS to start staking. Suppliers willing to accept more complex jobs immediately will be able to do so by staking their existing CUDOS from the outset.

Through this staking mechanism, suppliers who fail jobs can be penalised and suppliers who fake the outcome of a job to avoid expending computational work will lose their stake. In cases where jobs are failed due to non-malicious errors (e.g. connection failure), the stake will be held until that supplier completes another job. Upon completion of the next job, the conditions which led to the failure of the initial job will be assessed and a decision will be reached as to whether a penalty is required. If the job is completed successfully, the supplier will retain their staked tokens and the profits made for completing the job, along with a boost in reputation.

3.2.9 Reputation

Each supplier requires reputation to participate in the network. The more reputation possessed, the higher the priority granted to their networked devices and, by extension, the more jobs they will be eligible to receive.

Increased reputation may also improve network trust in a supplier. Reputation can be used by consumers as a metric to decide which machines will receive their jobs. Thus, reputation is a way for suppliers to ensure they receive the best jobs.

A tiered system will be introduced whereby suppliers are ranked by reputation. Entry-level consumers will start on tier 0, with some base reputation. This reputation can be increased by successfully accepting and completing jobs, and for successfully passing the tests sent. Other factors may come into play as well.

When a supplier has accumulated enough reputation, they will ascend to a higher tier, granting access to higher priority and better paid jobs (e.g. jobs requiring fast completion and which thus carry a premium).

Where a supplier has advanced security capacities, that supplier may be promoted instantly to an advanced tier. Security certificates (e.g. the SAS 70 standard for the US or the UK, or ISO 27001 for the EU) will ensure a higher tier for the suppliers holding it. Suppliers with secure hardware (e.g. Intel SGX or AMD SEV) will also be rewarded with increased reputation.

Various other factors may be taken into account in the reputation system, including:

- Success rate of test jobs
- Success rate of completed jobs
- Amount of time positively contributing to the network
- Number of devices provided
- Number of CUDOS tokens staked in the past

For more information on Cudo's reputation system, as well as hardware optimisation, benchmarking, auto-switching between cryptocurrency mining and compute jobs, encryption and use of blockchain, see the Cudo Technical Paper.

3.2.10 Projected Supplier Revenue

The baseline supplier revenues for running mining workloads are:

Standard desktop PC	\$1-30/month
Standard laptop	\$1-30/month
Gaming PC	\$15-150/month
Game console	\$10-50/month
Mining rig	\$150-1500/month
Mobile	< \$10/month

Cudo projects a 50-100% increase in revenues on average for compute workloads on a per hour comparison. A maximum 4x revenue increase is anticipated for a device at any time. Therefore, the projected revenues for compute workloads are:

Standard desktop PC	\$1.50-120/month
Standard laptop	\$1.50-120/month
Gaming PC	\$22.50-600/month
Game console	\$15-200/month
Mining rig	\$225-6000/month
Mobile	< \$15/month

3.2.11 Projected consumer savings

Consumers are calculated to save over 75% on average by switching from their current cloud provider to Cudo:

- 70% in CPU video rendering
- 75% in university research
- 85% in GPU video rendering
- Up to 95% in simulations

3.3 Target market

Cudo's target market can be split into three distinct groups. These are blockchain developers, supply side entities and demand side entities. Supply side entities include anyone capable of providing compute resources to the network. These can be data centres, miners, gamers or anyone with a computer, smartphone or similar device that is sitting idle. Demand side entities comprise anyone who requires computational resources or is willing to write software for the platform to facilitate the deployment of jobs. This is a far broader category since Cudo is able to service a wide range of computing demands. Examples include 3D rendering companies, entities requiring simulations, companies conducting medical imaging and developers of AI/Machine Learning tools.

3.3.1 Smart Contract Blockchains

Cudo offers a layer 2 compute network for smart contract platforms such as Ethereum. Being this layer 2 Turing-complete, Cudo will also play a key role in blockchains which are missing a Turing-complete layer to run computations, such as Algorand. Blockchain developers in any blockchain where the CUDOS network runs will be able to unlock new functionalities, offering for example escrow services or powering DeFi to the next level.

In addition, Cudo's vision is to provide blockchains the aforementioned layer 3, enabling any kind of workload on any kind of hardware to be requested from a blockchain directly. This will not only open new horizons to blockchain developers, but also to the wider compute community.

3.3.2 Cloud Supply Side

• HOSTING/SERVICE PROVIDERS

As well as consumer hardware, Cudo runs on data centre and hosting hardware. The data centre and web hosting market has a number of key players including the following entities.

Data centre providers: A number of large data centre providers including Equinix, Virtus and Digital Realty provide space in their data centres for other entities such as hosting providers to rent and fill with their own servers. Typically included in the price is permanent power, cooling and security. General businesses can use this for their own services while hosting providers can rent this capacity out to others that wish to use cloud services.

Hyperscale hosting providers: Large providers like Amazon, Google, Microsoft, Alibaba and IBM use huge amounts of capacity for their operations. To facilitate this, they can either build data centres of their own or rent data centre space on a large scale and fill it with servers. As previously noted, the top 24 hyperscale hosting providers spent \$75B on new servers and data centre space in 2017.

Hosting providers: The rest of the hosting market is made up of thousands of smaller companies like Digital Ocean, Peer1 and Pulsant with many millions of servers globally. These companies provide similar services to the large providers but often have a more personalised service or offer more niche/ optimised services. These providers range from over \$1B in revenue down to smaller providers in the \$100K's.

Each of the above providers often have substantial quantities of underutilised excess capacity. Since these entities are providing on-demand services, they always need to have unused capacity ready and waiting to provision out when requested. Though this spare capacity is paid for, it often remains unused, simply consuming electricity in order to remain continuously available in case of sudden demand. The value of the hardware continuously depreciates, whether it is used or not. Further, hosting providers are often committed to their data centre space, so their costs remain fixed.

Service providers can drastically reduce their idle compute and increase their profit margins by committing their spare capacity to Cudo. The management of this capacity will be the responsibility of the service provider. As an example, a service provider could retain a small percentage of spare capacity to meet client needs on demand, submitting the rest to cryptocurrency mining and compute jobs of various timescales.

Capacity used for cryptocurrency mining can be instantly redeemed without any negative consequences in terms of reputation or staking for the supplier. Capacity committed to compute jobs can be retrieved upon completion of jobs. Service providers can monitor client demand and, when their reserved spare capacity becomes depleted below a predetermined threshold, retrieve compute from Cudo and ensure that sufficient spare capacity is maintained. In this way, service providers can maintain high compute availability to meet customer demand while monetising that capacity when idle.

• GAMERS

There are 162 million game console owners in the US alone⁶⁴ and 49 million consoles are shipped worldwide per year.⁶⁵ In addition, there are over 1.3 billion PC gamers,⁶⁶ many of whom own specialist rigs that are well adapted to performing computation. Gamers are a tech-savvy demographic with advanced knowledge of hardware and software. A significant proportion of them are also familiar with crypto assets including tokens and have experience sending and receiving cryptocurrencies. As such, they are an ideal target market for Cudo, second only to miners. Users can select to receive payment in whatever form they prefer such as Steam vouchers or as credits for their preferred games platform

• MINERS

While other mining solutions may feature automatic coin selection, most of them neither adapt to the market nor make decisions to ensure the most profitable configurations are set at any point in time. Cudo's software, on the other hand, automatically changes overclocking and memory timing settings on the hardware depending on the mining or compute workload to maintain optimum performance from every piece of connected hardware. This results in significant improvements in profitability and reduces power consumption.

Through its existing Cudo platform, Cudo Ventures has cultivated an extensive user base of cryptocurrency miners. Cudo offers an alternative that can provide greater profitability. Miners will be given the option to automate switching between cryptocurrency mining and distributed computing. The software will monitor market conditions and benchmark the user's hardware to establish in real time whether mining or computing is more profitable, and select the appropriate option without the need for intervention by the hardware owner.

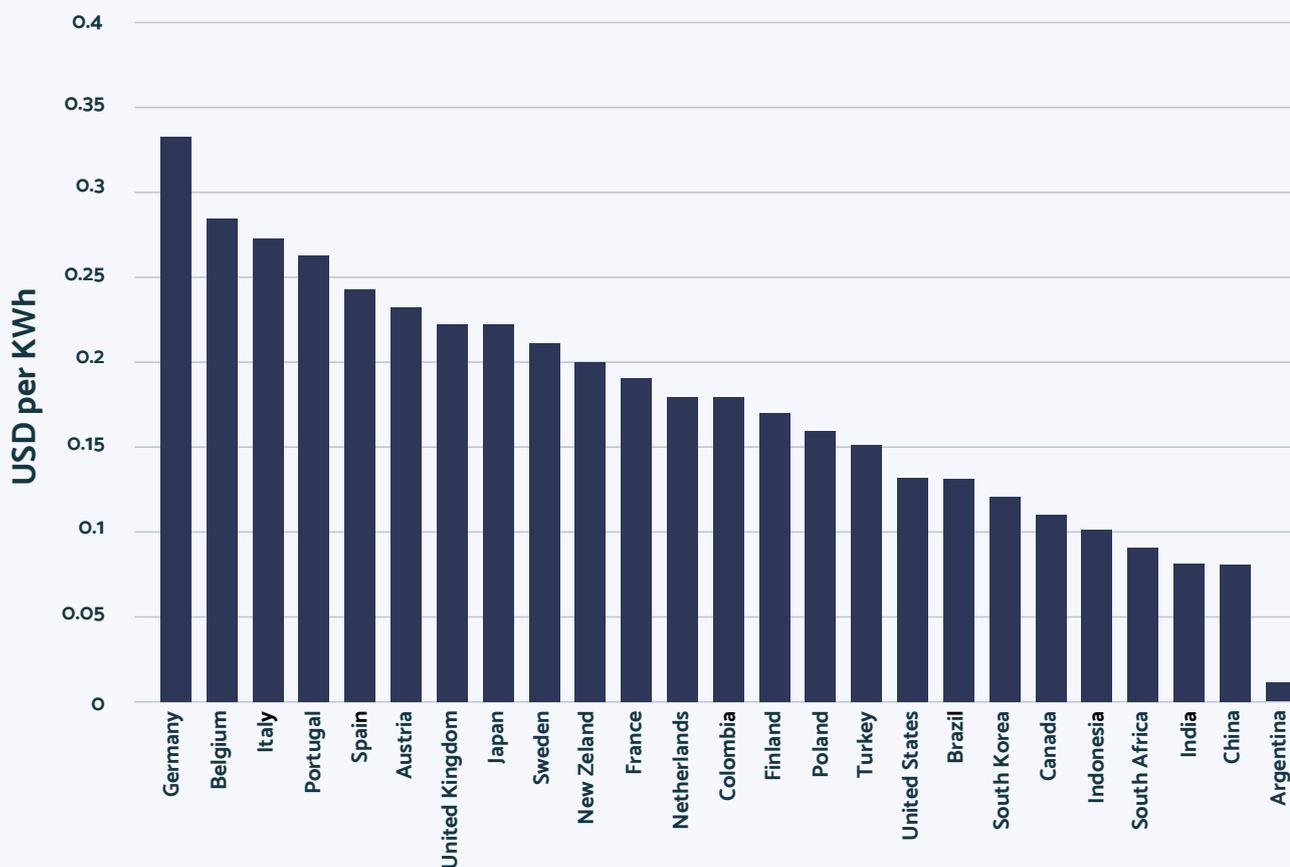
Cudo is suitable for casual cryptocurrency miners as well as professional mining operations, including ASIC farms. Professional mining operations are typically established close to cheap renewable energy sources such as hydro power. The flexibility of these operations enables them to locate where power is cheapest and cleanest. Adoption of Cudo by these players will provide an ecologically sound alternative to cloud data centres.

• **LOW COST ENERGY/LOW INCOME REGIONS**

In countries where energy prices are particularly cheap, Cudo will be more profitable and thus more attractive to suppliers. In developing countries, where there is a stable energy supply and access to desktop hardware, Cudo will be additionally valuable. This includes certain African and South American countries, where revenues earned from computing will go a lot further. Charities could gather old and unused devices and donate them to families in these regions to use to generate income. They could even assemble computing rigs made of old PCs or smartphones in these regions, assuming reliable power/infrastructure, and distribute the proceeds among the community.

• **THE UNBANKED**

Approximately 1.7 billion people in the world are unbanked with no identity, passports or digital money.⁶⁷ This precludes them from financial inclusion, including acquisition of cryptocurrencies, due to a lack of documentation. Anyone with access to a computer or smartphone, which a large percentage of the world’s population has in some form, can download Cudo and start to earn. Access to a steady stream of income can have profound effects on the lives of the unbanked. While this provision is possible with certain competitor solutions, Cudo breaks down the technical barriers to entry, allowing anyone to start earning money from spare computing capacity.



3.3.3 Cloud Demand Side

• RENDERING

3D rendering is used across a wide range of industries from product design to architecture and animation. Demand for high-quality, photorealistic rendering is increasing rapidly with the visualisation and 3D rendering software market set to progress at a CAGR of 22.5% between 2018 and 2025 to reach \$5.63B. However, modern rendering is incredibly GPU-intensive, and thus professional-grade rendering on regular computers is not an option. Depending on the specifications of a rendering job and the hardware used, a job can take a long time to complete and consume vast computing resources.

By distributing rendering jobs to hardware that is optimised for this task and by distributing large-scale rendering workloads among a number of devices within the network, Cudo will allow for the efficient and timely completion of jobs at low price.

• AI/MACHINE LEARNING

AI is a multifaceted tool that can include elements of language understanding, reasoning, learning and problem solving. Its applications encompass media, advertising, financial services, IT, telecom, retail, healthcare, transportation, and education among others. Spending on cognitive and AI systems is estimated to reach \$77.6B in 2022 (37.3% CAGR)⁶⁸ and the business value created by AI will reach \$3.9T in 2022.⁶⁹ Delivering AI solutions that are fit for purpose, however, is an incredibly costly, complex and time-consuming process.

Current cloud providers offer specialised, pre-trained AI solutions that greatly simplify AI-related tasks for consumers, but there are considerable costs involved that are beyond the reach of most individuals including researchers. Cudo offers an easy-to-use solution at a fraction of the cost.

• SIMULATIONS

Simulation software is in high demand across a number of market segments including medicine, the military, the automotive/aviation industries, and science.⁷⁰ Again, there is a vast and increasing demand for the computing required for the development of simulations from a variety of demand side players. The global simulation software market is anticipated to pass \$10B by 2025, with an 8.9% CAGR.⁷¹

Simulation software allows entities to simulate processes and analyse operations without physical execution. Simulations allow repeated testing with minor modifications, different inputs and other parameters at minimal cost when compared with physical execution. It also allows users to simulate events that could not be physically performed or analysed at scale (e.g. astrophysics modelling of neutron stars/black hole simulations). The creation of an accurate simulation requires significant computational resources and simulations often have to be repeated numerous times with a variety of inputs. The Cudo network can provide this computation reliably and at a low price.

• MEDICAL MODELLING/IMAGING

Medical modelling and imaging is becoming an essential tool for biomedical, life sciences and environmental research, as well as for everyday decision-making in health centres. Though cloud computing providers are currently offering medical imaging solutions, these services are not cheap. Even if they translate into cost reductions in the long run, many centres cannot afford the upfront costs.

In the medical industry, where lives are on the line, speed and accuracy are paramount. When a job has to be completed urgently, a priority premium can be paid to 'skip the queue' and receive the completed job back quickly, ensuring patients receive the correct treatment as early as possible (see section 3.2.5). With Cudo, consumers requesting the completion of medical imaging tasks can be confident in the accuracy of the returned jobs due to Cudo's reputation system (see section 3.2.9) and machine learning algorithms that detect fake work (see the Cudo Technical Paper). Cudo's advanced security and privacy encryption mechanisms will ensure that private and confidential patient information cannot be viewed or altered without permission. Cudo can also significantly lower the costs of these services, improve workflows and decision-making and reduce the need for costly and time-consuming administration.

4. CUDOS Token

A staking and compute token (CUDOS) has been engineered to power the CUDOS network as well as transfer part of the value generated on the Cudo network to users who stake CUDOS.

4.1 Token Utility

The (CUDOS) token is a qualification, discount and staking token. It is used to transact and stake in the CUDOS network and as a discount token to transfer part of the value generated on the Cudo network to users who stake CUDOS. CUDOS can also be kept, shared or donated to charity.

In the CUDOS network, the key functions of the token are:

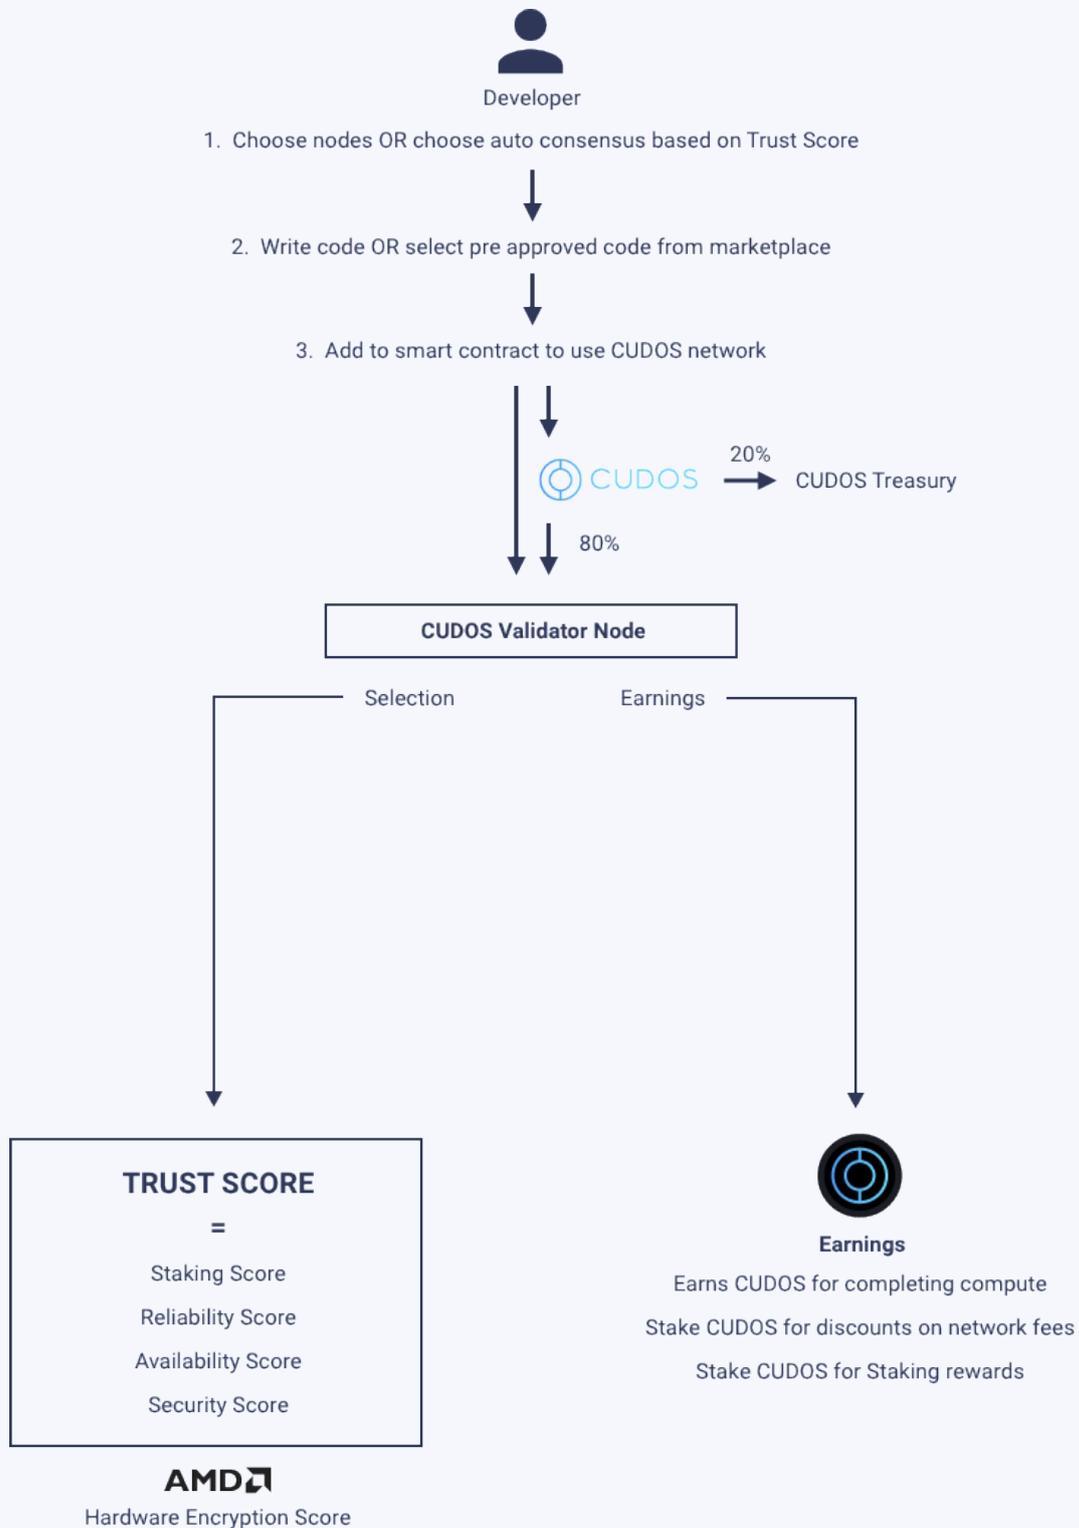
- 1, Staking 2,000,000 CUDOS to become a CUDOS Validator Node and get rewards
- 2, Delegated staking to support other CVNs
- 3, MoE powering the on-chain CUDOS network
- 4, Staking for discounts

In the Cudo ecosystem, the token can be used for:

- 1, Staking to receive a discount in fees or share the revenue from fees
- 2, Staking to qualify for jobs and earn greater revenue

4.2 CUDOS Network

The CUDOS network comprises a set of smart contracts allowing the off-load of on-chain compute work to off-chain resources via compute oracles. These compute oracles are the CVNs.



The network will initially be constrained to a maximum of 1,000 CVNs, belonging to at least 100 different hardware owners. As explained in section 3.1, blockchain developers can use the CUDOS smart contracts to select the CVNs where they would like their job to run. In turn, these CVNs will fetch the work, usually from an external data source, run the computation (and consensus when necessary) off-chain, and return the result.

The code run can either be written by the requesting developer, or can be chosen from a preselected template list for the most common workloads. This list will be found in the CUDOS marketplace.

Users willing to support an already existing CVN can do so delegating their CUDOS tokens. This is a non-custodial delegation, where users get rewarded for supporting the network. All staking rewards come from the CUDOS treasury, which collects a 20% fee of all compute tasks in the CUDOS network.

4.2.1 CUDOS Treasury

The CUDOS treasury objective is to collect part of the fees paid on the network for compute tasks, in order to support and incentivise the growth of the network and community. More precisely, the funds collected with the 20% fee, through slashing and through staking inefficiencies will be used to:

- Support the staking holder rewards
- Support development foundation
- Create development grants
- Incentivise and reward developers of marketplace templates
- Burning

Any leftover tokens from the treasury will be burnt. Burning is introduced in order to have a deflationary effect on the CUDOS economy.

4.3 Staking in the Cudo Platform

In addition to the usage just discussed for the on-chain CUDOS network, the token also has utility in the layer 3 platform, Cudo's compute layer, as mentioned at the beginning of the section.

Staking to receive a discount in fees

The CUDOS token utilises a discount token model based on a staking mechanism, where a fixed percentage of the total fee revenue generated on the Cudo network is distributed to users. This model is based on the fees paid by each user on the network and the amount of CUDOS each user has staked. In order to receive higher discounts, users must: 1) stake the relevant number of tokens required and 2) generate as much revenue as possible on the platform so as to increase the maximum discount available. The amount of revenue distributed to an individual user is capped and will never exceed the fee paid for using the platform.

The value of CUDOS is correlated to the activity of the Cudo network, based on the discount pool allocated to users, effectively aligning incentives between the owners of the project and holders of CUDOS. The growth/decline in the volume of computing power on the platform translates into higher/lower levels of discount allocated to users, which in turn drives the value of CUDOS. Rewarding users who hold on to CUDOS will eliminate the bootstrap problem and attract the necessary demand to meet the supply side, where the opportunity to reduce the fees paid to the platform by holding the token will positively impact user adoption.

There are two types of network participants:

Supplier: seller of compute resources who contributes their computing power to the networks. They effectively do the work.

Consumer: buyer of cloud services. People and companies who need computing power which they purchase from the network.

Cudo is the network in the middle which aggregates the suppliers and sells computing power to consumers. The network charges a percentage commission for each job completed on the platform.

The following formula is used to calculate the discount users will receive through the CUDOS token:

Define:

DP = discount pool, a portion of Cudo Ventures fee revenue contributed to discounts in a given period of time (every week, month etc).

R = discount rate, a percentage of revenues Cudo Ventures is contributing to DP, can be a fixed value, e.g. 50%, or a variable determined by a formula.

Cr = Cudo Ventures fee revenues in a given period.

DpT = discount per token, a value of discount users can enjoy for each token they hold.

N = multiplier, can be set by Cudo to increase discount per token in case the return on staking is deemed too low. Will always be set to a number $> = 1$.

TS = token supply, the number of tokens that participate in the distribution of discounts. The total circulating supply.

MiF = fees paid by an individual supplier "i".

TMi = token supply of an individual supplier "i" staked for discounts in the period.

MiD = discount enjoyed by an individual supplier "i".

Calculate:

$$DP = R * Cr$$

$$DpT = N * DP/TS$$

$$MiD = \min(DpT * TMi ; MiF * R)$$

As more/fewer players enter the ecosystem and want to lock more/fewer tokens, the number of tokens that these participants need to hold, in order to participate and benefit from the discount, adjusts accordingly.

This leaves enough tokens available to meet market demand while providing price stability as the circulating supply will adapt according to the performance of the network.

Supply side token flow:



Suppliers stake tokens to become eligible to receive compute jobs.



When a job is successfully completed suppliers get paid for the work.



When a job is not completed for non-malicious reasons (e.g. connection failures), the staked tokens are held until another job is completed.

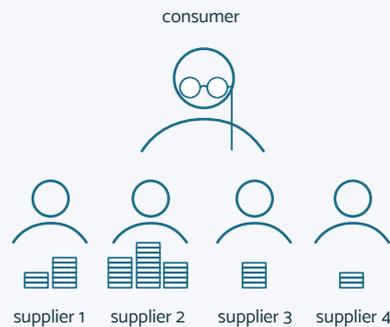


When a job fails for malicious reasons (e.g. faked work) the suppliers will lose their staked tokens.

Demand side token flow:



Consumers can pay for computing in a range of crypto assets and fiat.



Consumers can use the number of staked tokens as a metric for selecting a trustworthy supplier.

Staking for discounts:



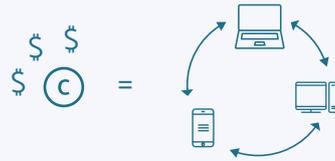
Suppliers can stake CUDOS to receive discounts on fees.



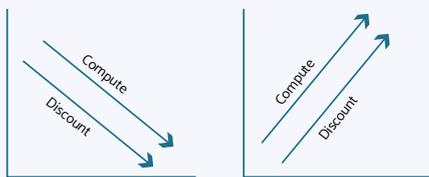
The greater the amount of CUDOS staked, the greater the discount received.



Assuming the discount rate is 50%, suppliers who stake enough tokens will see their fees reduced to half. The discount per user will be capped at the discount rate as well.



The value of CUDOS is correlated to the activity of the Cudo network.



Growth/decline in the volume of computing power on the platform translates into higher/lower levels of discounts allocated to users.



The absolute amount of discount provided by each token (and consequently its value) grows proportionally alongside network adoption.

Staking to qualify for jobs

The CUDOS token also utilises a staking/slashing mechanism, where in order to qualify for a given compute job, the supplier must have staked CUDOS tokens. CUDOS required will be proportional to each supplier's contribution to the network, with lower minimums for supplier's with more accrued reputation. This model serves to discourage malicious behaviour such as submitting fake work in order to game the system.

The slashing mechanism is executed when the supplier acts maliciously but not when a job fails due to non-malicious reasons such as a timeout (which may be due to a power cut or disconnected network for example). Initially, slashing policies and dispute resolution will be managed by CUDO's team directly. Following our path to decentralisation of the platform, once system dynamics and malicious behaviour is properly understood, slashing and potentially dispute resolution will migrate towards automated and community-driven policies.

4.4 Value Creation and Capture

Cudo creates value through a marketplace that monetises idle GPU and CPU capacity. CUDOS captures some of the value generated on the platform by entitling token holders to a certain level of discount made available on the network, assuming these entities contribute to the network and stake CUDOS.

This model provides clear incentives to hold CUDOS tokens over the long term, significantly reduces velocity and adds genuine utility. In time, the sum of cash flows/discount given to users should increase, helping to create a thriving ecosystem in which the goals of investors, asset providers, consumers and the Cudo team and advisors are aligned.

The CUDOS token will also increase user retention and lower user acquisition costs. In addition, it will incentivise users to generate more revenue for the platform by buying and providing more computing power and by recruiting other users.

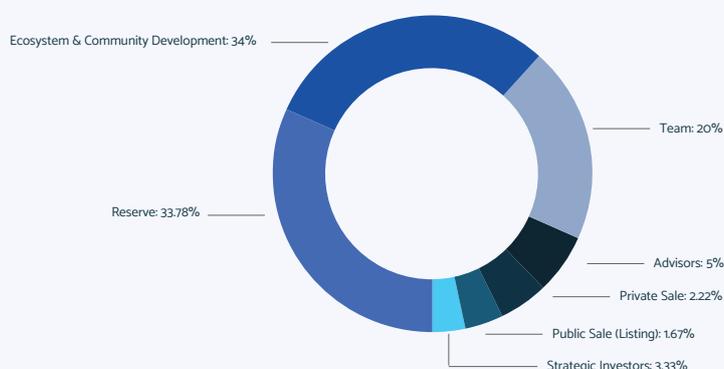
4.5 Token Sale

Private sale: 555,555,556 CUDOS	Public sale: 166,666,667 CUDOS
Total tokens for sale: 722,222,222 (7.22%)	
Hard cap: \$3,000,000	

Any unsold tokens in the sale will remain for at most 12 months after TGE with the issuer, and will be released periodically.

4.6 Token Distribution

The total token supply is 10,000,000,000 CUDOS, which are distributed and used as follows:



Allocation Bucket	Purpose
Private Sale and Strategic Investors	Tokens sold to seed investors and VCs and accredited investors during private rounds. Any unsold tokens from the token sale will be transferred to Ecosystem and Community Development fund.
Public sale	Tokens for exchange listing or launchpad platform.
Advisors	Tokens for advisors who assist with platform creation and sales.
Ecosystem & Community Development	To build the community and ecosystem via rewards, prizes, bounties, developer programmes, partnerships and grants.
Reserve	Used to provide additional funds for administration, operations, marketing and R&D or any other spend if needed.
Team	Founders and incentives for the core team.

4.7 Token Schedules

Private Investors

Early-stage investors receive a bonus as a reward for supporting Cudo's next stage, with different vesting terms and no lock up period. Strategic round investors comprise the first \$1m and buy the CUDOS token at \$0.003/CUDOS with 12 months vesting. The private investor round accounts for the second \$1m raised, with a token price of \$0.0045/CUDOS and 6 months vesting. Token release and accessibility will be daily, with 1/365 and 1/181 tokens released daily respectively.

Advisors

Tokens allocated to Cudo's advisors will be locked for 6 months after the Token Generation Event (TGE). Following this will be a vesting period of two years with tokens released each quarter.

Team

Tokens allocated to the team will be locked for 6 months after TGE. Following this will be a vesting period of three years with quarterly releases.

Reserve

The reserve will vest over a 120 month vesting period with quarterly releases (no lockup).

The release will follow a mathematical formula built into the vesting contract:

The token release ramps down with an average of 1.56% tokens of the Reserve released quarterly in the first 9 months and an average of 5.24% tokens of the Reserve released in the following 9 months. The remaining allocations will be gradually released at a decreasing rate after 18 months over the following 8 years and 6 months. All Reserve tokens will be released over 10 years from the TGE.

Ecosystem & Community Development

Tokens will be released over a 120 month vesting period with monthly releases (no lockup).

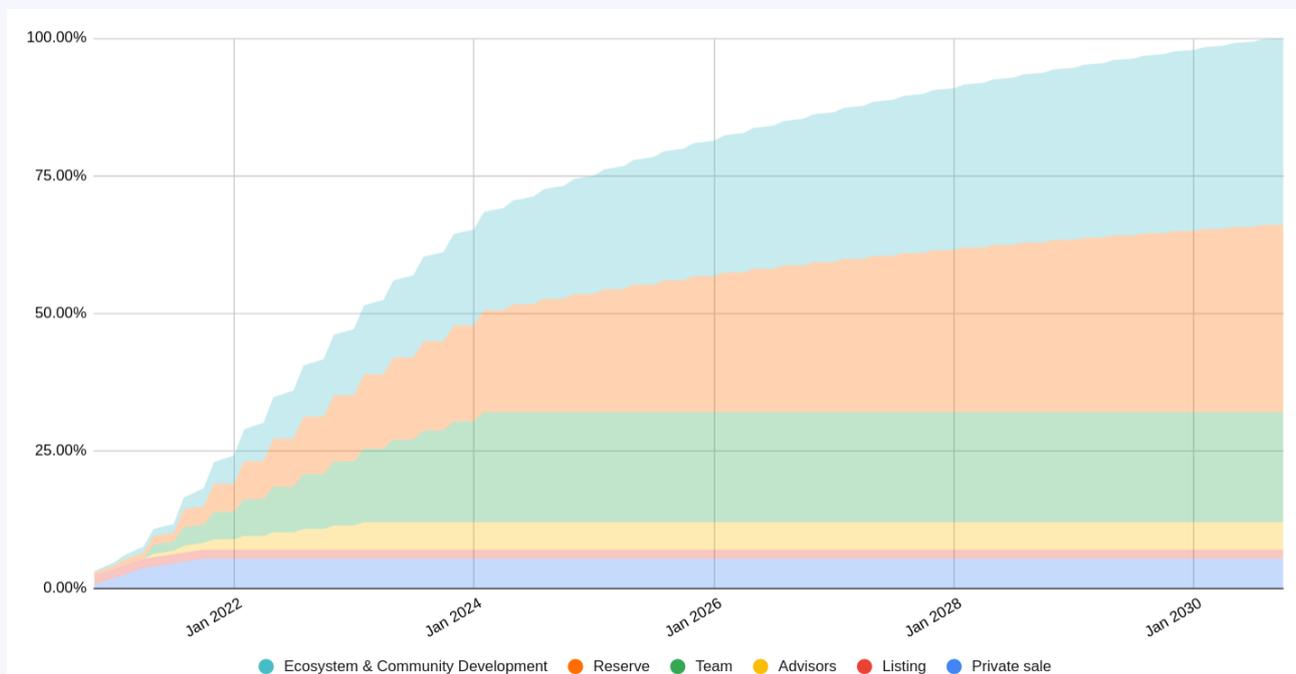
The release follows a mathematical formula built into the vesting contract:

The token release ramps down with an average of 0.52% tokens of Ecosystem & Community Development released monthly in the first 9 months and an average of 1.73% tokens of Ecosystem & Community Development released in the following 7 months. The remaining allocations will be gradually released at a decreasing rate after 16 months over the following 8 years and 8 months. All Ecosystem & Community Development tokens will be released over 10 years from the TGE.

The Ecosystem & Community Development tokens will be transferred to company wallets once a month.

Public Sale

Public sale tokens will be available from day 1 after TGE.



4.8 Use of Funds

Funds raised from the CUDOS token sale will be allocated transparently, through auditable transactions, supported by regular reports available to the public:

• ADMINISTRATION AND OPERATIONS (10%)

Used to cover administrative and operational costs. These include office rent, insurance, legal costs, office supplies and professional fees.

• RESEARCH AND PRODUCT DEVELOPMENT (30%)

30% of the funds raised will be dedicated to research and product development. This budget is key to building a robust computing platform powered by the CUDOS token. Research and product development activities include but are not limited to: compute workload research, smart contract development, platform development, API development, Cudo marketplace development, Cudo client development, testing/QA, bug fixes, deployment, release management, hardware optimisation and website development.

• **MARKETING (20%)**

To create awareness of the token and the ecosystem to help its adoption and continued growth, CUDOS will spend 20% of the funds raised on marketing activities such as social media impact, marketing campaigns and PR. This budget is essential in order to keep the community up to date with the latest platform developments, as well as to create awareness and educate about CUDOS’ objectives and milestones.

• **NODE VALIDATORS (10%)**

In order to cultivate and grow the node validators ecosystem during the early and mid-stages of the network, CUDOS will spend 10% of the funds to ensure that all node validators are rewarded for the contributions to the network. This budget is necessary to reach maturity of the overall ecosystem without any need to compromise the balance and adoption of the platform.

• **USER ACQUISITION (13%)**

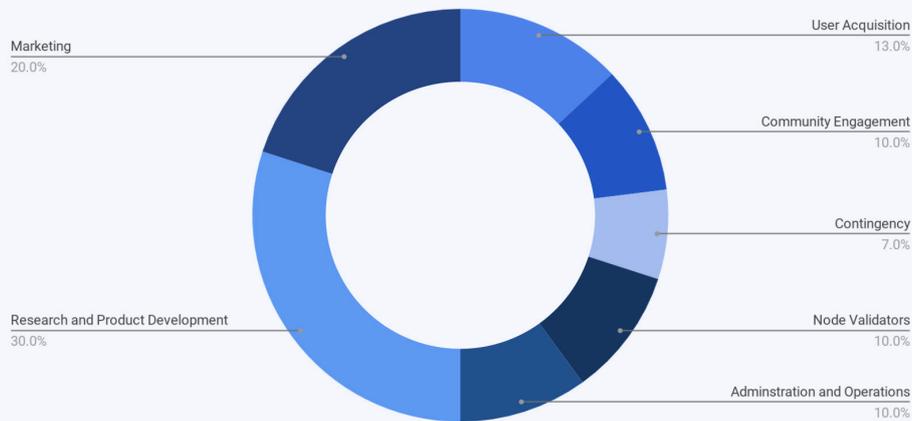
To grow the user base on both the supply and demand side, CUDOS will spend 13% of the funds raised on user acquisition activities such as marketing campaigns, referral programmes, advertising, PR and partner development. This budget is necessary to create a network effect to increase the value provided to suppliers and other participants of the network.

• **COMMUNITY ENGAGEMENT (10%)**

Community engagement is crucial to the success of token-based platforms. CUDOS will appoint dedicated community managers to educate, support and grow the community of suppliers, developers, users and token holders. The budget will also be spent on organizing developer meet-ups and events to foster relationships between members of the CUDOS community.

• **CONTINGENCY (7%)**

The contingency budget will be held for future unforeseen expenditure.



4.9 Security

Cudo Ventures is partnering with accredited companies to implement best practices for all security and regulatory processes. A leading smart contract auditor has been appointed to ensure the code governing the CUDOS token is free of bugs and works as intended.

Treasury and future allocated tokens will be stored with a trusted and insured third party. This will provide oversight for token distribution, and ensure secure storage for all tokens and capital.

5. Use Cases

Cudo's technology is suitable for most batch-based, scaleout or parallel workloads. The platform will launch with video rendering initially on the cloud side, but there is a wide range of use cases that external developers can create applications for.

5.1 Blockchain Compute

Blockchains missing a Turing-complete layer 2 cannot offer such as escrow or advanced payment systems. Turing-complete blockchains are lacking mass adoption due to the lack of flexibility, as running complex compute jobs with external data is either inviable or not possible at the moment. Similarly, while DeFi growth has now stalled due to gas costs⁷², further services may require advanced techniques such as Montecarlo simulations, which will need off-chain support to run effectively.

5.2 Rendering

Cudo's compute platform lowers the barriers to creating high quality CGI, providing rendering tools to make the process easier and utilising the power of the thousands of GPUs in its distributed computing network to quickly render complex CGI images for a fraction of the price of cloud rendering. For GPU rendering, by using existing GPU cryptocurrency mining rigs as well as gaming PC's, the distributed computing network will have the capacity to render videos on demand while rewarding GPU owners for the hardware provided. For CPU rendering, gaming PC's and service providers' spare server capacity will provide the necessary infrastructure. As a result, industry-leading companies as well as individual consumers will benefit from a cost-effective solution to video rendering while having the opportunity to contribute to the network and earn revenue when their own high-powered devices are sitting idle.

5.3 AI/Machine Learning

Consumers will be able to run the models they require by accessing computing power to process custom-made algorithms and creating and partitioning data to run in template models provided by the network's embedded AI solutions. Cudo will also integrate tools including TensorFlow and Scikit-learn to facilitate model-building within the network.

5.4 Simulations

Cudo can offer extremely efficient and cost effective simulation computation. As with rendering, Cudo's compute platform will select devices on its network that are best suited to meet the demands of simulations and perform the required computations efficiently and cheaply.

Where, for example, a car manufacturer wishes to perform crash simulations or aerodynamics modelling to ensure that the product is as safe and efficient as possible, thousands of minor alterations may be made to a host of variable inputs from road conditions to the shape of the chassis, speed and wind resistance. When on a budget and a tight release schedule, it is important that these simulations can be completed quickly and cheaply. Given that the design decisions made in light of the outcome of these simulations could directly impact the safety of those using the final product (i.e. drivers and passengers), it is vital that the simulations are accurate.

By submitting these jobs to the Cudo network, the car manufacturer can be sure that the completed simulations will be returned quickly and that the results are accurate due to Cudo's staking and reputation systems as well as its fake-work detection mechanisms (see Cudo's Technical Paper).

6. Business Development

6.1 Current Cudo Platform

Prior to the development of the layer 2 CUDOS network and the layer 3 Cudo compute platform, Cudo has created and launched a professional mining application, for GPUs, CPUs and ASICs. While the early focus was on crypto enthusiasts and charities, the company has now widened its range to support professional mining operations, in the form of mining farms, data centres and businesses.

Cudo allows anyone with a PC, dedicated mining rigs or ASICs to instantly begin mining cryptocurrency, and to efficiently manage and monitor all mining hardware. The service launched publicly in 2019 and users are onboarding rapidly. There are currently 150,000+ signed up users and over 20,000 devices (nodes) running the software at any time. A major contributing factor in Cudo's success is its ease of use. Users can "plug and play" before leaving the built-in mining profitability algorithm to take care of the rest. Cudo monitors the cryptocurrency mining market, including prices and the mining difficulty of each coin, in real time. It analyses the data collected and cross-references it with data on each mining device in the network to select the most profitable settings and coins for each piece of hardware. Cudo's automatic coin switching and overclocking algorithms ensure maximum efficiency and profitability at all times. This removes the need for manual system configuration, one of the most difficult and time-consuming elements of entering into cryptocurrency mining.

Similarly, professional mining operations benefit from Cudo's personalisation and close support. Cudo is constantly pushing new features into the platform, to improve and refine its monitoring and integration of large scale mining operations.

For more advanced users, Cudo Miner offers total configurability. Miners are free to bypass Cudo's automation and manually set up their operations, selecting their preferred coin to mine and managing their overclocking settings. Cudo's platform is available in multiple languages, is supported for Windows, Mac and Linux and will eventually integrate smartphone support.

Cudo also allows users to donate their computer's idle resources to various charities. At the moment, Cudo allows cryptocurrency miners to contribute the coins that they have mined, automatically converting them to the preferred payout currency of the selected charity and delivering those proceeds. The option to donate profits will be made available to all of Cudo's supported workloads when they come online.

Another focus on the market for Cudo are gamers, as Cudo allows users to generate value from the underutilised computing capacities of their gaming hardware during off-time or while in play.

Cudo is currently developing whitelabels of its software for existing game industry platforms and brands. By integrating Cudo's software, these entities can enable their users to generate credit that can be spent within the platform ecosystem. This presents an innovative new way for platforms, game publishers, developers and distributors to monetise without increasing costs for users. It also provides a way for platforms to engage with users who do not traditionally make in-game purchases.

6.1.1 Cudo's Ecosystem

Cudo is a growing business, taking the next step into its development. It already has a large community, and an extensive and broad pre-existing client base with exceptional hardware diversity. As such, Cudo has been setting the grounds for this moment for the last three years, and has gathered the necessary experience, knowledge and equipment to push forward all its general-purpose scalable compute solutions.

7. Market Space

7.1 Summary

Taking a broad overview of the various market players, both centralised and decentralised, several trends become apparent. On the CUDOS side, Ethereum 2.0 is progressing very slowly, and the first layer 2 alternatives are starting to launch their mainnets, attracting a lot of interest from the community. On the more traditional compute Cudo side, Centralised cloud computing is already well established and has the first mover advantage. It also has the support and backing of some of the biggest players in the business and computing world. However, centralised computing has several disadvantages as addressed earlier, particularly in regards to security, service outages, costs and energy consumption.

Grid computing promises to overcome many of these issues, but decentralised cloud computing projects have been unable to gain traction for a number of reasons which Cudo has addressed:

- **Sound token economics**

Each of the decentralised competitors cited in this section have implemented hybrid payment tokens that can be replaced by any other form of payment. Those tokens are required to use the network and therefore introduce a barrier to entry. In addition, they lack fundamental value, encourage speculation leading to volatility and create no customer loyalty. As a result, they fail to capture value. In contrast, Cudo has developed a discount token, which creates a good value proposition for token holders as its utility price is connected to the success of the platform. By doing so, Cudo has avoided pricing jobs in a potentially volatile currency at the same time ensuring a consistent billing method.

- **Ease of use**

The current crop of decentralised computing products are complex and highly technical in nature. For this reason, the solutions are not designed with the 'everyman' or 'everywoman' in mind. This creates barriers to entry, hampering their ability to gain future market traction. Cudo's products have a user-friendly UI which does not require any technical expertise to set up and run.

- **Guaranteed 24/7 user revenue**

Users of other decentralised platforms rely on the availability of compute jobs. If there are not enough jobs available then some suppliers will be left with no revenue streams. Cudo offers 24/7 revenue, as crypto mining is an optional workload which is always available.

- **Customer base**

Most other decentralised projects have a clear lack of supply and demand in the platform, and thus the amount of nodes and transaction volume for these projects has not scaled. Cudo already has around 20,000 concurrent devices at any time, which ensures a head start over the decentralised competition, and is also constantly seeking and closing partnerships with leading companies in the hardware, cloud and blockchain space to ensure enough compute demand.

- **Pragmatism over ideology**

Cudo strongly supports the concept of decentralisation, but believes that this can only be achieved over time. Rather than attempt to have the entire system decentralised on day one, a process of “rolling decentralisation” makes more sense while customers and market traction is gained and usability is improved.

- **Specialised experience**

Cudo has years of experience developing apps and software that distribute jobs to workers, as well as data centre expertise. The Cudo team has also carefully analysed the market and all the competition, and learned, planned and addressed the main challenges that arise. This research, together with Cudo’s experience, ensures a direct path towards a solid, working product.

Cudo’s package is attractive to enterprises accustomed to working with leading cloud providers, while incorporating elements of the Web 3.0 stack including tokenisation and distributed data storage. As such, Cudo has a number of advantages over centralised cloud computing, including:

- **Low-latency services**

Cloud providers such as GCP and AWS cannot guarantee low-latency computing, as they have a limited number of locations and users need to physically connect in a building in order to get a low latency connection. This is not suitable for many applications or on-site hosting, as it is very costly, adds a barrier, takes time to arrange and also it generally means users need to host it locally. With Cudo users can choose where the network runs, and due to the decentralised, edge computing nature of the platform users can select virtually any location in the world to be less than 4ms from the data or compute location. No other hyperscale providers are able to offer this at the moment.

- **Environmentally friendly**

As outlined above, ecologically speaking current cloud technologies are not sustainable. By utilising already existing hardware and maximising its computing time, Cudo offers a much needed sustainable solution to the world’s increasing computing needs.

7.2 Layer 2/Oracle Projects

Recently many projects have started providing blockchains with oracle solutions, in order to allow smart contracts to fetch external data. The most popular current use case is DeFi, and most projects use PoS and delegation to incentivise the ecosystem and validate the data provided.

Chainlink

Chainlink is a project providing cryptocurrency prices to smart contracts. They provide decentralised oracle networks which have been widely adopted by Ethereum’s DeFi community, taking extra care in the quality of the data. While an early implementation suffered from high gas costs, the Chainlink team is working on alternatives to overcome its issues.

DIAdata

DIA, which stands for Decentralised Information Asset, is an oracle platform providing DeFi markets with financial data. It also aims to provide traditional financial data, to make the market more accessible and transparent and avoid single points of failure for critical data points, upon which many financial products are based.

Band Protocol

Band Protocol is a data oracle platform that aggregates and connects real-world data and APIs to smart contracts. The oracle network runs on its own Tendermint-based blockchain using the Cosmos SDK. Band not only focuses on data for DeFi applications, but also on other kinds of real-world data such as sports or weather.

Matic Network

Matic is trying to solve blockchain's scalability issues by using sidechains for off-chain computation, via the Plasma framework and a PoS network with validators. As such, it offers a high transaction throughput in sidechains, with finality achieved on the mainnet of the layer 1 chain it is built upon. Its first compatible blockchain is Ethereum.

NEAR Protocol

NEAR is a layer 1 protocol which uses sharding and PoS to provide scalability and low-cost transactions. Mainnet was launched in May 2020, and as of September 2020 the ecosystem is fully run by the community. As of the time of writing, there are dozens of independent nodes running the network.

Algorand

The Algorand Protocol is a PoS blockchain using a novel multi-signature scheme that achieves substantial savings in bandwidth, storage requirements, and verification effort. The focus is on very low transaction fees and atomic transfers, in order to allow frictionless finance.

Cosmos

Cosmos is an ecosystem of blockchains that can scale and interoperate with each other. These blockchains are tied to the Cosmos Hub, a PoS blockchain that uses the ATOM cryptocurrency. However, each blockchain is independent and secures itself. In addition to blockchain interoperability, Cosmos focuses on scalability and usability, with a modular framework for developers that allows fast blockchain building.

Polkadot

Polkadot is a blockchain protocol that unites a network of child blockchains (parachains), in order to allow them to operate together and scale. The parent chain is called the Relay Chain, using a consensus algorithm code-named GRANDPA. Every user holding Polkadot's native currency, DOT tokens, is able to participate in the network's governance and council election.

7.3 Other Blockchain Projects

iExec

A multifaceted project that centres around blockchain-based decentralised cloud computing. iExec is building a P2P marketplace for computing resources, integrates dApps and has a strong team, many of whom are greatly respected in the fields of grid computing and blockchain. iExec is a member of the OpenFog Consortium which includes ARM, Cisco, Dell, Intel, Microsoft, and Princeton University.

While its RLC token was initially used for payments only, it also has staking utilities now. However, pricing and payments in the platform are still done in RLC, creating inherent price volatility, and the install and set up process requires technical expertise.

iExec has partnered with Intel and IBM for job security, but the platform has not yet been widely adopted by the community. The main challenges the project faces are:

1. Mass user adoption (both from the supply and demand sides).
2. No announced strategic partnerships to increase user base.
3. Pricing and payments are volatile due to their nature.
4. Decentralisation from the start slows down user adaptation and tech delivery.

Cudo has studied the project carefully and developed a strategy which avoids the aforementioned challenges, by:

1. Having an initial user base of over 20,000 monthly active users.
2. Actively seeking partnerships and signing deals to ensure demand and supply customers.
3. Pricing jobs in fiat currency to ensure price predictability and stability.
4. Keeping some centralised components at the start built for easy decentralisation.

Golem

Golem is a generalized distributed compute platform that focuses on developer experience. Originally developed with an initial focus on CGI rendering, Golem released Brass Golem version 0.22 on January 14 2020, considerably behind the original roadmap. After that, the company decided to depart from the original whitepaper idea, and have started working on a new version of the platform.

DFINITY

DFINITY describes itself as “a new blockchain computer that is similar in concept to Ethereum but has vastly improved performance and, ultimately, unlimited capacity.” DFINITY is a not for profit organisation, and thus is not a market competitor for Cudo.

Following a long-term research project approach, DFINITY launched in 2019 Mokoto, a programming language to be used in their internet computer technology. However, as of the time of writing no sample code is available for the general public.

Filecoin

Filecoin offers a distributed storage solution, allowing users to share and use spare storage. It is currently in testnet, and at the end of September 2020 started its phased approach to launching mainnet.

Filecoin focuses only on distributed storage, and thus it is not a direct competitor for Cudo. Cudo is expecting to launch similar decentralised storage services further down the road (please see section 8 for the roadmap), once the main compute platform is up and running. Therefore, the development efforts of both companies are complementary at the moment.

7.4 Centralised Cloud Computing

Centralisation of services has many downsides, as all entities need to trust service providers and any policies they decide to implement. The cloud computing industry is currently centralised and monopolised by a few big players, which makes security and trust a central issue. These big companies are also constantly building new infrastructure to meet the increasing demand, which is having a massive ecological impact on the planet and also translates to expensive services. As described above, Cudo naturally addresses these issues by setting a path to decentralisation and using already existing hardware to provide cloud services.

• Amazon Web Services

For the last decade or so, AWS has led the cloud IaaS market.⁷³ It provides a wide range of cloud-based products and services including compute, storage, databases, analytics, networking, mobile, developer tools, management tools, IoT, security and enterprise applications.⁷⁴ AWS generated revenues of \$25.6B in 2018 and \$35B in 2019.⁷⁵ AWS provides clients with a vast tool set that continues to grow, but the service has been known to suffer outages and other technical issues.⁷⁶

In addition, AWS' centralisation means that users are geographically-limited when choosing where they want their compute jobs to run. At the time of writing, AWS has 24 locations,⁷⁷ which greatly limits data sovereignty and high latency connectivity for most of the world. Furthermore, users need to pay the ability to scale, which makes these services very expensive. Typically this means up to 10x the cost on Cudo.

• Google Cloud

Google Cloud specialises in high compute offerings like big data, analytics and machine learning and offers considerable scale and load balancing.⁷⁸ Google uses deep discounts and flexible contracts to try to win projects from clients that are currently spending significant sums with cloud competitors.⁷⁹

However, Google Cloud locates also in a very limited number of countries.⁸⁰ Just like AWS this is not only limiting data sovereignty and high latency connectivity for most of the world, but it also increases the price of the compute up to 10x with respect to Cudo's offering.

• Microsoft Azure

Microsoft has an elite enterprise background and decades of experience. Azure is particularly versatile due to the ability to integrate with Windows and other Microsoft products and there are various discounts available to loyal Microsoft users. However, Azure customers have reported issues with technical support, documentation and training. Finally, limited locations dampen it in terms of data sovereignty and high latency offering.⁸¹

• Alibaba Cloud

Alibaba Cloud is the biggest cloud services provider in China, with the greatest coverage within the country. However, outside of Mainland China it only has 12 locations.⁸² This leads to all the same issues regarding data sovereignty, high latency and high prices as described above.

8. Roadmap



9. Team

9.1 Core Team



Matt Hawkins / CEO & FOUNDER

Founder of numerous UK tech businesses including C4L, a co-location, connectivity, cloud and communications provider which provided around 1% of the UK's internet infrastructure and won fast growth awards from Times Tech Track 100, Deloitte UK Top 50, and Deloitte EMEA Fast 500. Matt ran C4L for 16 years, going on to sell it in 2016 in order to finance Cudo Ventures. He was involved in M&A during the transition period, helping acquire other businesses. Matt has been involved in mentoring young entrepreneurs in the UK, is an Entrepreneur in Residence at Bournemouth University and was awarded Entrepreneur Of The Year at Dorset Business Awards and HSBC's South West Business Thinking Winner.



Andrew Sturme / ENGINEERING DIRECTOR

Andrew has 19 years of industry experience in software development and solutions and is adept at designing lean, highly-supportable and reproducible systems in hosted and server-provider environments. Prior to joining the Cudo team, Andrew was the lead architect at C4L. Andrew's extensive experience in building platforms that support hundreds of thousands of users and that require the highest levels of uptime and resilience gives him the ability to rationally and logically reason complex problems and to develop solutions.



Andrew Walker / CHIEF REVENUE OFFICER

Over 15 years experience in SaaS growing & leading cross-functional teams to build world class products that deliver customer success. Chief Revenue Office @ Cudo and advising Startup Founders & Leadership teams of Pre-seed to Series B/C on their journey by helping them to understand, plan & achieve their targets. Andrew previously co-founded @Clicktools (acquired by Survey Monkey & Callidus Inc.)



Lee Woodham / CHIEF OPERATING OFFICER

A strategic leader with a proven track record of driving business growth resulting in increased profit margins and revenue in international multi- faceted organisations. Transforming business through change management, the introduction of innovation and technological support and rigorous corporate governance, in global blue-chip and Private Equity backed start-up organisations.



Nuno Pereira / VP OF PARTNERSHIPS

Nuno was Group Sales Director for Foresolutions, responsible for the sales and marketing of the three companies within the group. A telecommunications business, after 13 years with DSD Food Group managing over 100 sales reps. He is experienced in building B2C and B2B relationships and creating strong strategic relations.



Pete Hill / VP OF SALES

Pete previously worked for C4L where he built and managed a 250-strong partner channel that included over 30% of the top 100 managed service providers and value added resellers. With a BA (Hons) degree in Computing & Internet Technology and 12 years IT industry experience, Pete has a vast technical understanding of cloud and internet service provider technologies.



Peter Willis / TEAM LEAD

With over twenty years of experience in web and mobile software development as a Co-Founder and Director of Sycora, Peter is an expert in architecting, designing and developing platforms and services that require access to millions of data points, specialising in software process improvement and software engineering innovation. Peter has a BA (Hons) degree in Computing & Internet Technology from Bournemouth University.



Mia Whitewood / SR BUSINESS DEVELOPMENT ENTERPRISE

Mia joins Cudo with over 23 years in IT & Cloud across EMEA & ASIA PAC and will help drive further success and growth at Cudo. Mia excels in delivering strategic business value and has driven major profitable growth at start-ups and global players including; egg.com, Pipex, UUNET, Rackspace UK and Salesforce.com & RedEye Apps in Australia.



Sean Berry / HARDWARE ARCHITECT

With a BEng (Hons) in Medical Electronic Design, Sean has experience in solving problems in fields of scientific computing such as system dynamics and fluid dynamics. He is skilled in developing digital signal processing algorithms and embedded systems. Sean has worked on a variety of medical engineering problems and medical systems as well as working with technologies such as radar and sonar.



Dr Joan Garcia Tomo / DATA SCIENTIST

Holds a PhD in Theoretical and Mathematical Physics from the University of Southampton. Prior to joining Cudo, Joan worked primarily in education at the Universitat de Barcelona and the University of Southampton.



Victor Leach / SOFTWARE ENGINEER

Left Bournemouth University with a First Class Honours degree in Games Programming. Since then Victor has gained experience in a number of software developer roles as well as directing his own limited company, Vmlweb, offering various development services.



Joe White / SENIOR SOFTWARE ENGINEER

Studied Digital Media Design at Bournemouth University with a specialisation in Web Development. Gained experience as an intern web developer at Zeta Agency before taking on a role at Mad River Ltd. as a junior web developer.



Dexter Edwards / LEAD SOFTWARE ENGINEER

Having studied Mathematics to Masters' level (MMath), specialising in Finance for Masters year project, Dexter won multiple awards at University of Reading, and achieved a First Class Honours degree. Dexter previously worked on asset finance calculation engines for large clients including BMW, Alpha Romeo, Toyota and Xerox. Other experience includes development work at world-leading financial derivatives and online trading companies.



Ashley Smith / CLIENT JOURNEY ADVOCATE

Ashley has significant post-university experience in marketing, acquiring a level 4 certificate in digital marketing from the Chartered Institute of Marketing. Before joining the Cudo team, Ashley gained industry experience as a Digital Marketing Apprentice with IDE Group.



Chris Saganic / LEAD SOFTWARE ENGINEER

Chris has designed, developed and maintained websites, web apps and native mobile apps for UK and major international companies such as Navachi, South Devon College, Ahoy.club, BMW & Volvo. He is a full stack developer and UI designer, specialising in Vue & Node.js. Chris has been involved in 22 projects in the last 3 years since graduating with a BA (Hons) degree in Digital Media Design from Falmouth University.



Pedro Luz / SENIOR SOFTWARE ENGINEER

A highly focused software engineer with +10 years experience in a variety of development and engineering positions at companies including Moola, onebillion, Bluebell and SPORT BUFF. Methodical and a keen eye for detail results in solid coding and trustworthy software programmes



Avi Mutham / LEAD SITE RELIABILITY ENGINEER

Cloud Native DevOps and SRE engineer with an accomplished track record of a software technologist and consultant with 12 years of industry experience working across multiple business domains which includes Banking/finance, E-Commerce, Retail, Travel and logistics.



Arthur Verrept / JUNIOR SOFTWARE ENGINEER

Currently studying BSc (Hons) Internet Design at Plymouth University, In his first year and second year achieved First-Class Honours and made the dean's list for the 2018 year placing him in the top 12% of students in his respective faculty. Arthur is currently doing a year long placement with Cudo Ventures where he works on creating features for the website and maintaining both front and back end code.



George Lezeu / JUNIOR FRONTEND DEVELOPER

After graduating from Birmingham City University with a First-Class Honours degree in Criminology, Policing and Investigations, George decided to change careers, recently joining the software industry. Translating his skillset in problem solving and strong attention to details into a more technological approach, George joined Cudo as a Junior Frontend Developer eager to further his skills in Vue.js.



Richard Poole / LEAD DEVELOPER

With over 19 years of experience in software engineering and web development, Sycora co-founder Richard is skilled in front and back-end development and full stack development and design. Richard has a BSc (Hons) degree in Computing from Bournemouth University.



Dominic Burns / PROJECT MANAGER

Starting out in defence telecommunications and building on his engineering skill set with companies such as Marconi Communications upon leaving the military. Dominic has spent the last 15 years managing successful IT industry sector projects, ranging from cloud, network and data centre services with Telecity Group and C4L. To leading cross functional project teams in the delivery of efficiency programmes, driving continuous improvement and the provision of exceptional service.



Neethu Stephen / FINANCE

Neethu is a finance professional with over 10 years of post-qualification experience. She has worked across the globe and her prior experience includes both scale-ups as well as MNCs like Ernst&Young and Komatsu.



Jo Watson / HUMAN RESOURCES

Over 18 years experience, I've managed and supported global software projects, business change, transition of services and HR programmes, as part of global transformation.

Jo has a strong background in HR, frequently managing IT organisation change projects, which I feel has grounded me with many transferable skills such as strong communication and relationship/stakeholder management.



Sonja Greenfield / EXECUTIVE ASSISTANT

With over 17 years experience providing support to executive teams, Sonja has worked across a range of industries including magazine & online publishing, finance, education & health. Prior to joining Cudo Sonja worked for a health tech company, assisting with scaling the business & launch of their SaaS products into the UK market.



James Dickie / GRAPHIC DESIGNER

Since graduating from Bournemouth University with a BA (Hons) degree in Industrial Design, James has gained just short of 10 years design experience and over 5 years marketing experience with companies such as Animal Friends Insurance, BlueBee Solutions and Greenspan Projects. James is always keeping a close eye on advances in technologies within the design industry which lower carbon footprints and can help save the environment.



Flavius Berciu / CLIENT JOURNEY ADVOCATE

When living in 3 different countries I became aware of differences in people's culture. It is a great experience for acquiring a global perspective, a concept at the base of every successful company. I have always been interested in customer service and that can be seen in my work experience from serving a drink at the bar, assisting the elderly with daily activities, to assisting customers with a product issue. I always aim to provide empathy with my service.



Jeffrey Elliott / COMMUNITY MANAGER

After studying Law at BPP University Jeffrey went on to build his own educational investment website and community, from there Jeffrey went on to begin working with blockchain companies helping them to build their own communities from scratch and raise funds during their token sales.

9.2 Advisors



Chris Deering / FORMER PRESIDENT OF SONY COMPUTER ENTERTAINMENT EUROPE

Chris Deering is an American businessman and marketer best known for his role as president of Sony Computer Entertainment Europe. He is credited as one of the main marketers of the successful PlayStation and PlayStation 2 video game consoles earning him the title of 'The Father of the Playstation'.



Jörg Roskowitz / DIRECTOR OF PRODUCT MANAGEMENT & BLOCKCHAIN TECHNOLOGY, RADEON TECHNOLOGY GROUP AT AMD

As Director of Product Management & Blockchain Technology for the Radeon Technology Group at AMD, Jörg is the Blockchain Technology lead working with strategic Enterprise partners on Blockchain based solutions.



Sean Li / CO-FOUNDER & CEO, FORTMATIC

Sean is the co-founder and CEO of Fortmatic, an alternative to MetaMask without needing any installations. Prior, he was the creator of Kitematic, a tool making Docker accessible on Mac and Windows. After its acquisition by Docker, he became the product lead for Docker Desktop touching the lives of millions of developers every month.



Chris Gale / CO-FOUNDER, VERASITY

A serial entrepreneur with a passion for digital and advertising technologies. Co Founder at Veracity.io, an attention based platform for video rewards which raised \$18m in an ICO in 2018. Founder of an early rich media mobile ad platform which sold to Phunware Inc. which went public reaching a market cap of \$2.6B. He is also an active investor and advisor in web 3.0 businesses.



Malcolm Tuck / MANAGING DIRECTOR, ESET UK

Malcolm is an experienced senior executive and current Chairman of JAM Technologies Limited. Previously a Board member and Managing Director of Kaspersky Lab the worlds 3rd largest, IT Security vendor. Malcolm joined Kaspersky in 2008 as UK MD and achieved significant market share over an 8 years period.



Craig Fletcher / GAMER & TECH EVANGELIST, ESPORTS EXPERT, FOUNDER OF MULTIPLAY & INSOMNIA GAMING FESTIVAL

Craig founded Multiplay in 1997 and led its growth to become the UK's largest video gaming events and Esports company, as well as the largest provider of online gaming servers - acquired by GAME Digital PLC in 2015. Craig then led the strategy and partnerships around Esports and competitive gaming for the GAME Group, until leaving in 2017. He now acts as a business angel, consultant and investor in the Esports, media, tech and video gaming space.



Maggie Fang / FOUNDER & CEO, WYNSUN CAPITAL MANAGEMEN

Founder & CEO of Wynsun Capital Management, a PE fund whose portfolio includes Alibaba, Uber and WeWork. Maggie also is an adviser to early and growth stage tech companies in the fields of social media, AI, blockchain, etc. Previously, Maggie worked in the prestigious Banks Goldman Sachs and Morgan Stanley.



Jay Coshan / FOUNDER & CEO, UNIQUE BLOCK GROUP

Specialising in Investment and strategy on cryptocurrency, blockchain and the UK property sector, Jay founded Unique Block Group, which invests in tokenized projects related to blockchain technology that are creating the new decentralised economy. Jay invested in over 30 blockchain focused start-ups across the globe. He also continues to oversee multiple property development and investment projects.



Duncan Cook / FOUNDER, 3 SIDED CUBE

Founded award-winning app development company 3 Sided Cube in 2009 whose clients include Lloyds, Red Cross and RNLI. Duncan's apps have been recognised for saving lives, receiving over 10 million downloads.



Joshua Riddett / FOUNDER OF EASY CRYPTO HUNTER

Josh is an award winning entrepreneur and a leading expert in the UK cryptocurrency scene. Founder of crypto service provider Easy Crypto Hunter, Josh has keynoted at major blockchain and cryptocurrency conferences as well as consulting and advising for some of the biggest names in the industry and mentoring new companies entering the space.



Philip Forte / BLOCKVENTURE COALITION

Partner at the BlockVenture Coalition (the largest alliance of university blockchain groups & VCs in the world). Founder & former President of the Blockchain Group at Carnegie Mellon University. Previously an Institutional Asset Management Analyst at PNC.



Tyler Wellener / PARTNER BLOCKVENTURE COALITION

Tyler is a partner at the BlockVenture Coalition (the world's largest alliance of university blockchain groups & VC funds). He has experience working as an Associate at Struck Capital Crypto and a Consultant at Deloitte.



Jonas Sevel Karlberg / CEO & FOUNDER, AMAZIX

Jonas is the founder of AmaZix and the co-founder of the Nordic Blockchain Association. He has more than 20 years of commercial experience working in retail operations management, business IT systems development and franchise management for some of the world's largest brands. Working with the Nordic Blockchain Association, Jonas is involved with building and nurturing a blockchain community within Scandinavian countries, providing startups with a tangible environment for interaction and networking. As the Managing Director of AmaZix, Jonas leads a global team of 60, specializing in blockchain advisory and community management.



Olena Clayton / EXECUTIVE COACHING

Experienced at building and running world class teams for startups and enterprises including Barclays Investment Bank, Olena's true passion lies in executive coaching for entrepreneurs and their businesses. With degrees in computer science and finance and qualifications in NLP and coaching, Olena brings enormous value and experience to Cudo.



Andy Gray / CO-FOUNDER OF BLOCKROCKET.TECH

Specialises in Blockchain engineering and consulting. Advocate of decentralisation.

9.3 Partners

OUTLIER VENTURES

Outlier Ventures is a venture firm that invests in companies building Web 3.0 with decentralised technologies. They invest in Convergence-related startups. Convergence sees blockchains and other decentralisation tools like smart contracts, oracles, and tokens, as new decentralised infrastructure enabling technologies such as AI, IoT, 3D Printing, Robotics and Mixed Reality to converge. Cudo joined OV's accelerator programme in December 2019..

AMAZIX

AmaZix is the world's largest blockchain advisory firm with four divisions comprising Digital Advisory, PR & Marketing, Legal and Corporate Finance. It has worked with more than 120 projects that have raised a total of \$1.3 billion, including HDAC, Bancor and Bankex. The AmaZix Digital Advisory team has designed the token economics of projects issuing a wide range of utility and security tokens, optimising these for value capture and incentive alignment to satisfy all network participants.

EASY CRYPTO HUNTER

The UK's premier cryptocurrency service provider, specialising in crypto mining rigs, blockchain and ICO consultancy and cryptocurrency accounting.

GAMEKIT

Gamekit is a gaming reward platform that gives gamers gift cards and discounts every time they play games and complete quests. Gamekit is the biggest platform of its kind in the world, with over 16 million users.

BLOCKVENTURE COALITION

The BlockVenture Coalition is a global alliance of top research groups and VC funds in the blockchain space. The firm has partnered with 41 university blockchain research groups and 44 blockchain-specific VCs. BVC is looking to grow the broader distributed ledger space by helping universities collaborate and nurturing top talent as well as providing resources to the best startups.

OTOY

OTOY Inc. is the definitive cloud graphics company, pioneering technology that is redefining content creation and delivery for media and entertainment organizations around the world. OTOY's award-winning technology is used by leading visual effects studios, artists, animators, designers, architects and engineers, providing unprecedented creative freedom, new levels of realism, and new economics in content creation and distribution powered by the cloud.

DISTRIBUTED COMPUTE LABS

Distributed Compute Labs is a Canadian educational non-profit responsible for developing an ECMAScript-based distributed compute platform to accelerate Canadian science, innovation and discovery. The Distributed Compute Protocol (DCP) is a framework for distributing computational workload over a network of heterogeneous nodes. DCP enables arbitrary IoT devices and digital infrastructure - anything from smartphones to enterprise web-servers - to contribute idle CPU and GPU cycles to configurable and collaborative compute networks.

10. Closing Statement

We are reaching a peak in computing efficiency. Bell's law posits that a new computer class emerges every 10-15 years to accommodate demand. For the last decade, centralised cloud computing has dominated, but this system is unsuited to the needs and nature of tomorrow's networked devices.

Cudo is spearheading the next phase in the evolution of computing. This era will be centered around the utilisation of spare capacity and distributed networks, enabling entirely new industries with new peripherals and applications to emerge.

Blockchain technology is a remarkable candidate to lead the way to this next paradigm, and Cudo will be at the forefront of the technological leap necessary for it to reach mass adoption. Cudo offers scalable and decentralised data and compute solutions, which are the main issue hindering blockchain's real-world applications. This enables more efficient DeFi and increased integration with data feeds and internet facing API's.

Furthermore, the distributed grid computing infrastructure we've developed is sustainable at scale from the traditional cloud perspective. Cudo offers an alternative to ecologically destructive computing practices by reducing the creation of additional hardware and data centres and enables enterprises to increase productivity while reducing costs.

I've been fortunate to witness the rise of cloud computing, and to have helped businesses thrive through effective utilisation of cloud-based services. Now, as the industry prepares to enter its next phase, I look forward to introducing enterprises to the many benefits distributed computing has to offer. As a crypto and blockchain enthusiast myself, combining cloud computing expertise with the recent blockchain developments has been a really exciting challenge for myself and the team which is now becoming a reality.

The Cudo team is committed to realising this vision and I'm confident that solutions like Cudo's compute platform and the CUDOS network will serve a vital role in Web 3.0, helping to foster a fairer and more efficient internet where privacy and open access are guaranteed.



Matt Hawkins
CEO and Founder

11. Glossary

Blockchain: a distributed and decentralised public ledger recording all transactions for a given cryptocurrency.

Cloud Computing: the use of networked facilities for the storage and processing of data rather than a user's local computer, access to data or services typically being via the internet.

Consumer: One who requires the completion of computational work such as an SME, enterprise, service provider, NGO or researcher.

DeFi: stands for Decentralised Finance, an ecosystem of financial decentralised applications running on blockchains.

Distributed Cloud Computing: the interaction of cloud technologies across multiple geographic locations to perform computations.

Edge Computing: the processing of data near to the network's "edge" - the place where the data is generated - rather than transmitting that data to a data centre for processing. This processing can be done by the device producing the data or by a local computer or server.

Fog Computing: the midway point between cloud and edge computing, with data processed at multiple points within the network between the edge and the cloud.

Grid Computing: the practice of using numerous widely separated computers (esp. ones linked via the internet) to carry out large computational tasks by sharing processing power, typically for the purposes of scientific research.

IaaS: Infrastructure as a Service, an instant computing infrastructure, provisioned and managed over the internet. A vendor provides clients pay-as-you-go access to storage, networking, servers and other computing resources in the cloud.

Layer 1: also referred to as the main chain, it is a blockchain system such as Bitcoin or Ethereum, with typically low-transaction throughput.

Layer 2: a secondary network built on top of the underlying layer 1 blockchain, typically solving layer 1 issues such as scalability. Loopring, which uses ZK-Rollups, is an example of a layer 2 network for Ethereum.

Mesh: a network in which a number of computers or processors are connected together, specifically a mode of connection in which each computer or processor is connected to a number of others to form a multidimensional lattice.

PaaS: Platform as a Service, in which a service provider offers access to a cloud-based environment where users can build and deliver applications. The provider supplies underlying infrastructure. In addition to storage and other computing resources, users are able to use a suite of prebuilt tools to develop, customise and test their own applications.

Provider: One in possession of computational resources who offers those resources to the Cudo network for the purpose of completing jobs put forward by consumers.

SaaS: Software as a Service, a software distribution model in which a third-party provider hosts applications and makes them available to customers over the Internet. SaaS is one of three main categories of cloud computing, alongside infrastructure as a service (IaaS) and platform as a service (PaaS).

12. Endnotes

1. <https://www.symantec.com/content/dam/symantec/docs/white-papers/data-privacy-and-compliance-in-the-cloud-en.pdf>
 2. https://www.greenpeace.org/usa/wp-content/uploads/2019/02/Greenpeace-Click-Clean-Virginia-2019.pdf?_ga=2.261944584.481306289.1561645043-1903329051.1561645043
 3. <https://www.youtube.com/watch?v=cidZRD3NzHg&feature=youtu.be>
 4. <https://www.cambridge.org/core/books/computing-universe/lickliders-intergalactic-computer-network/OEC06B278505D-9671DF38398C725E3B1>
 5. <https://www.statista.com/statistics/505251/worldwide-infrastructure-as-a-service-revenue/>
 6. <https://www.forbes.com/sites/louiscolombus/2019/04/07/public-cloud-soaring-to-331b-by-2022-according-to-gartner/#257a6ef45739>
 7. <https://www.srgresearch.com/articles/hyperscale-data-center-count-jumps-430-mark-us-still-accounts-40>
 8. <https://www.cloudcomputing-news.net/news/2018/mar/01/new-figures-show-staggering-capex-levels-hyperscale-cloud-providers/>
 9. <http://telecoms.com/495537/google-makes-13-billion-worth-of-cloud-plans-for-2019/>
 10. <https://www2.deloitte.com/us/en/pages/consulting/articles/a-catalyst-for-cloud-the-flexibility-to-weather-the-covid-19-related-downturn.html>
 11. <https://www.thoughtworks.com/insights/blog/macro-trends-tech-industry-may-2020>
 12. <https://www.statista.com/statistics/830485/world-fog-computing-revenue-by-vertical/>
 13. <https://www.statista.com/statistics/1101442/iot-number-of-connected-devices-worldwide/>
 14. <https://a16z.com/2016/12/16/the-end-of-cloud-computing/>
 15. <https://bitcoin.org/bitcoin.pdf>
 16. <https://www.investopedia.com/terms/s/satoshi-nakamoto.asp>
 17. <https://www.cudominer.com/introduction-to-blockchain-and-cryptocurrency-mining/>
 18. <https://www.blockchain.com/charts/hash-rate>
 19. <https://blockgeeks.com/guides/ethereum-gas/>
 20. <https://www.forbes.com/sites/niallmccarthy/2019/07/08/bitcoin-devours-more-electricity-than-switzerland-infographic/#e7fb89921c0e>
 21. https://near.org/wp-content/uploads/2019/10/Economics_in_Sharded_Blockchain.pdf
 22. <https://www.cudominer.com/blockchain-related-papers/>
 23. <https://coinmarketcap.com/currencies/bitcoin/>
 24. <https://www.bbc.co.uk/news/technology-42237162>
 25. <https://www.cityam.com/what-is-decentralized-finance-defi/>
 26. <https://defipulse.com/>
 27. <https://cointelegraph.com/news/waiting-for-ethereum-20-but-layer-two-solutions-are-needed-now>
 28. <https://www.businessinsider.com/companies-waste-62-billion-on-the-cloud-by-paying-for-storage-they-dont-need-according-to-a-report-2017-11>
 29. <https://www.businessinsider.com/companies-waste-62-billion-on-the-cloud-by-paying-for-storage-they-dont-need-according-to-a-report-2017-11>
 30. <https://etherscan.io/chart/transactionfee>
 31. <https://a16z.com/2020/04/30/crypto-fund-ii/>
 32. <https://www.nextplatform.com/2017/02/24/glimmer-light-dark-silicon/>
 33. <https://www.sciencedirect.com/science/article/pii/S0065245818300147>
 34. <https://www.micron.com/about/blog/2018/october/metamorphosis-of-an-industry-part-two-moores-law>
- Fixed broadband services with a download speed of at least 30 Mbit/s.
Connections that can deliver download speeds of up to 1 Gbit/s.
35. https://www.ofcom.org.uk/_data/assets/pdf_file/0020/130736/Connected-Nations-2018-main-report.pdf
 36. <http://unohrlls.org/custom-content/uploads/2018/01/Worlds-least-developed-countries-on-target-to-achieve-universal-and-affordable-internet-by-2020.pdf>
 37. <https://www.srgresearch.com/articles/hyperscale-data-center-count-jumps-430-mark-us-still-accounts-40>
 38. <http://www.clickclean.org/downloads/ClickClean2016%20HiRes.pdf>

39. <https://www.wired.com/2006/10/cloudware/>
40. <http://www.clickclean.org/downloads/ClickClean2016%20HiRes.pdf>
41. <https://www.wired.com/2006/10/cloudware/>
42. <http://www.clickclean.org/downloads/ClickClean2016%20HiRes.pdf>
43. <https://www.cloudcomputing-news.net/news/2018/apr/10/google-achieves-100-renewable-energy-target-becoming-first-public-cloud-do-so/>
44. <http://www.clickclean.org/downloads/ClickClean2016%20HiRes.pdf>
45. <https://www.scientificamerican.com/article/cloud-computings-substantial-footprint/?redirect=1>
46. <https://www.theguardian.com/environment/2015/sep/25/server-data-centre-emissions-air-travel-web-google-facebook-greenhouse-gas>
47. <https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle>
48. https://www.insight.com/content/dam/insight/en_US/pdfs/apc/apc-estimating-data-centers-carbon-footprint.pdf
Edinburgh Centre for Carbon Management Ltd., Forestry Commission Scotland Greenhouse Gas Emissions Comparison Carbon Benefits of Timber in Construction, August 2006
49. https://www.afcom.com/Public/Member_Resources/Learning_Center/Research/Public/Resource_Center/Researchs_Public.aspx?hkey=812b2bf5-a7bf-4fb8-9d5e-17bc206d979c
50. <https://www.forrester.com/report/Worldwide+PC+Adoption+Forecast+2007+To+2015/-/E-RES42496>
51. <https://www.statista.com/statistics/330695/number-of-smartphone-users-worldwide/>
52. <https://www.dualshockers.com/ps4-92-million-lifetime-sales-2018/>
53. <https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/>
54. https://webcache.googleusercontent.com/search?q=cache:IELO7IH6_UJ:https://www.forbes.com/sites/benkepes/2015/06/03/30-of-servers-are-sitting-comatose-according-to-research/+&cd=1&hl=en&ct=clnk&gl=uk
55. <https://aws.amazon.com/blogs/aws/cloud-computing-server-utilization-the-environment/>
56. <https://www.independent.co.uk/news/uk/home-news/office-workers-screen-headaches-a8459896.html>
57. https://eta.lbl.gov/sites/default/files/publications/computers_lbnl_report_v4.pdf
58. <https://assets.kpmg/content/dam/kpmg/pdf/2015/11/cloud-economics.pdf>
59. <https://www.businessinsider.com/companies-waste-62-billion-on-the-cloud-by-paying-for-storage-they-dont-need-according-to-a-report-2017-11>
60. <https://hal.archives-ouvertes.fr/hal-02083080/document>
61. <https://ieeexplore.ieee.org/document/7439752>
62. <https://www2.deloitte.com/content/dam/Deloitte/at/Documents/human-capital/at-2018-deloitte-human-capital-trends.pdf>
63. <https://www.marketingcharts.com/cross-media-and-traditional/videogames-traditional-and-cross-channel-82362>
64. <https://www.statista.com/statistics/276768/global-unit-sales-of-video-game-consoles/>
65. <https://www.statista.com/statistics/748072/number-pc-gamers-world-platform/>
66. <https://globalindex.worldbank.org>
67. <https://www.idc.com/getdoc.jsp?containerId=prUS44911419>
68. <https://www.gartner.com/en/documents/3868267/forecast-the-business-value-of-artificial-intelligence-w0>
69. <https://www.millioninsights.com/industry-reports/simulation-analysis-software-market>
70. <https://www.alliedmarketresearch.com/simulation-and-analysis-software-market>
71. <https://cointelegraph.com/news/defi-mainstreaming-impossible-until-dexs-integrate-layer-2-experts-say>
72. <https://www.statista.com/chart/18819/worldwide-market-share-of-leading-cloud-infrastructure-service-providers/>
73. <https://aws.amazon.com>
74. <https://www.statista.com/statistics/233725/development-of-amazon-web-services-revenue/>
75. <https://downdetector.com/status/aws-amazon-web-services>
76. <https://aws.amazon.com/about-aws/global-infrastructure/>
77. <https://www.datamation.com/cloud-computing/aws-vs-azure-vs-google-cloud-comparison.html>
78. <https://cloud.google.com/compute/vm-instance-pricing>
79. <https://cloud.google.com/about/locations>
80. <https://www.gartner.com/doc/reprints?id=1-2G2O5FC&ct=150519&st=sb&alid=1154870580>
81. <https://azure.microsoft.com/en-gb/global-infrastructure/geographies/>
82. <https://www.alibabacloud.com/global-locations>

