TATAU

Blockchain-based, enterprise-focused, distributed supercomputing platform

Whitepaper
October, 2018

tatau.io

Abstract

Artificial Intelligence (AI) and predictive analytics are accelerating the rate of discovery in science, helping facilitate new medical breakthroughs, and bringing new technologies and tools to life. AI is a disruptive force in all industries for the greater good. Advancements in AI technology, like deep learning, have changed everything from the information on a Facebook news feed to how medical therapies are used to treat illness. AI is changing our very basic daily interaction with the world from phones, transportation, news to health and will continue to do so at an accelerated pace. Even with all of these current applications of AI, there are many more that will be realized in the years to come.

The world will be AI driven within the next few years. It's an exciting world, and it requires vast amounts of hardware and computational power.

Without accelerated and scalable advanced computational capacity, AI will not be able to deliver on its promises. Currently, huge investments are being made in data centers, which utilize a traditional general-purpose central processing unit (CPU) platform.

However, building more data centers is not enough to meet the growing demand in a world with an insatiable appetite for AI.

AI's growing demand for computational power can be effectively met by creating a solution that utilizes both massively parallel graphics processing unit (GPU)-based computing for increased speed and less energy consumption and distributed processing, using available computers in a sustainable manner that reduces the environmental footprint caused by increased computing.

Tatau.io is an enterprise-focused, blockchain-based distributed computing platform established to meet the growing demand for computational power required for complex computations. Our initial focus will be in verticals such as AI and deep learning, but will ultimately include areas such as video rendering and predictive analytics. Tatau's computational supply harnesses existing distributed GPU-based computing resources owned by crypto-mining and gaming companies. Tatau is able to pay these companies more for their compute power than they can earn mining or providing computing power for gaming, while at the same time provide the computing resource to AI companies for less than they pay to traditional compute suppliers.

Tatau leverages blockchain implementations including Ethereum, BigchainDB and IPFS to provide a high level of security in computation, information, and payment.

The Tatau token provides two uses in the Tatau ecosystem. Firstly, the token is used as payment, including buying and selling computational output. Secondly, the token is used for staking computational nodes and validators in the network.

Today, we are seeing a world where machines and technology function in ways that were once unimaginable. The team at Tatau embraces the possibilities inherent in the growth of AI with a keen focus on creating a best-of-breed platform that will excite and inspire its clientele.

The abstract is an introduction to the white paper. Any decision to invest in or purchase the virtual financial assets should be based on consideration of the white paper as a whole by the investor. The offering of virtual financial assets does not constitute an offer or solicitation to sell financial instruments. Any such offer or solicitation of financial instruments will be made only by means of a prospectus or other offering documentation in terms of any applicable Maltese law.

Table of Contents

Background

Today, AI powers our phones, computers, household thermostats, voice-activated virtual assistants, and our online purchases. It's just the beginning of the transformative impact that artificial intelligence will have on society. By acting like a capital-labor hybrid, artificial intelligence offers the ability to amplify and transcend the current capacity of capital and labor to propel economic growth.[1] We see a future where AI will be used to manage the use of scarce resources, assist in scientific exploration, accelerate the diagnosis and delivery of healthcare, and venture into a range of new realms that are beyond our imagination.

As artificial intelligence expands into different fields, there is a continual push to analyze more data on more complex models in order to gain better predictive insights. The massive amounts of data used in AI, video rendering, and predictive analytics require equally massive amounts of computational power. This computational power comes in large part from cloud services offered by renting out data server space, with the revenues for these services falling in the billions.[2]

The growth in AI has traditional CPU-based data centers scrambling to keep up with the increasing demand from new more computationally expensive algorithms like deep learning.

As a result, investment in new data centers is booming with a record year of investment in 2017 and expected growth of between 12-14% over the next two to five years.[3] To meet the growing demand for computational capacity data centers are increasingly becoming GPU-based computing data centers. In fact, in the first six months of 2017, Nvidia tripled its data center revenue.[4]

[1] https://www.accenture.com/us-en/insight-artificial-intelligence-future-growth
[2] https://seekingalpha.com/article/4084040-large-deep-learning-data-center-market?page=2
[3] https://www.forbes.com/sites/bisnow/2017/09/29/double-digit-growth-expected-in-data-center-industry/#63964589 44ea
[4] https://seekingalpha.com/article/4084040-large-deep-learning-data-center-market

The recent interest in GPUs comes from the unique benefits that GPUs offer when applied to artificial intelligence and graphics rendering. Since GPUs can process 100s of times more threads than current CPUs, they can be used to parallelize many of the operations needed in artificial intelligence and video rendering. This means that GPUs allow for faster calculations and have more memory bandwidth, and they can process data at greater speeds than CPUs. To top it all off GPUs can deliver all this while using one tenth the energy required by CPUs.

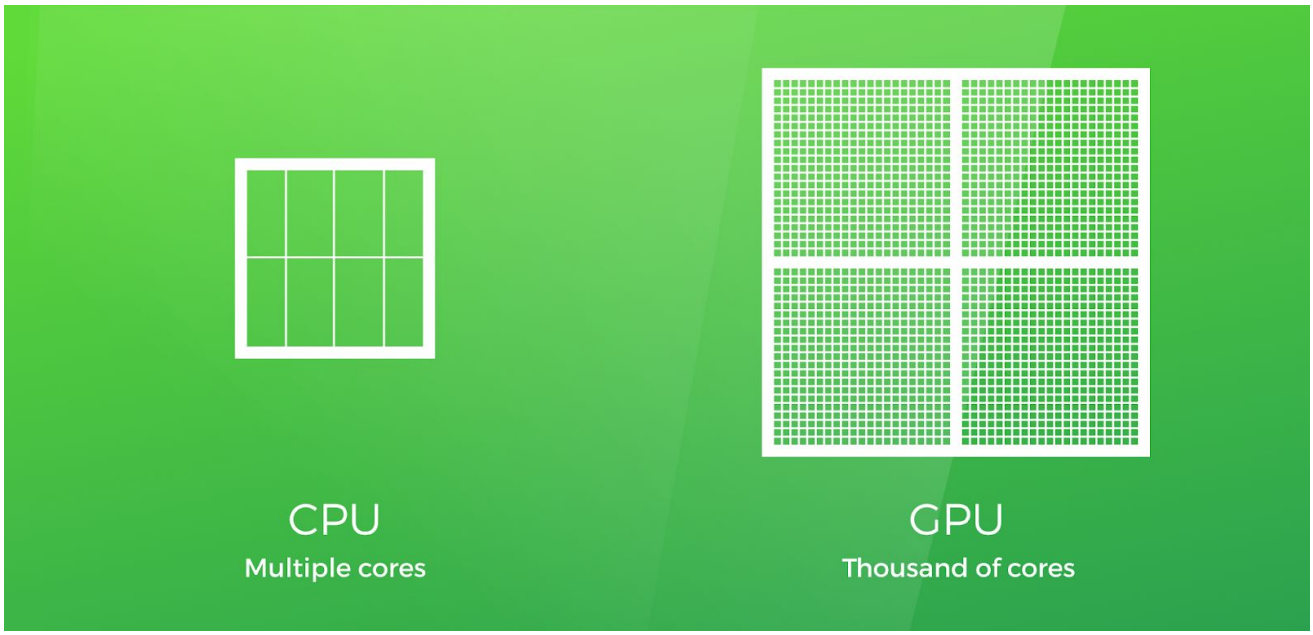

Problem

Building more traditional CPU-based data centers is not going to be enough to meet the rising growth in AI. The CPUs in these centers are inadequate to cope with the exponential growth of artificial intelligence, video rendering and predictive analytics. Furthermore, these data centers are some of the most energy-intensive building types, consuming 10 to 50 times the energy per floor space of a typical commercial office building. Collectively, these spaces account for approximately 2% of the total U.S. electricity use.[5]

It is not a question of capacity – with billions of consumer computers that could be utilized – the worldwide computing capacity has yet to be fully realized. And, we have not yet fully utilized GPU power.

For AI, predictive analytics, and video rendering to scale, we must find an effective way to boost cloud services by harnessing the vast computational capacity of GPU powered computers in a secure, eco-friendly, and efficient way.

Solution

The issue of growing demand for computational power by AI can be effectively met by creating a solution that utilizes both GPU-based computing and distributed processing for increased speed and lower energy consumption. The Tatau platform uses available computers in a sustainable manner that reduces the environmental footprint caused by increased computing.

Tatau is a system that offers a solution to AI's ever-increasing demand for computational capacity. The Tatau Platform is a blockchain-based platform that offers an advanced, eco-friendly cloud service specifically designed for processing AI and video rendering utilizing GPUs. By incorporating blockchain components such as Ethereum in its stack, Tatau enables anyone to buy and sell computing capacity on its decentralized computing platform. This promises to take advantage of new advancements in GPU-based efficiency algorithms for artificial intelligence and 3D rendering, and fill the need for more eco-friendly computational solutions. Furthermore, by utilizing the deep neural net training algorithm proposed in Training Deep Nets with Sublinear Memory Cost, Supplier Nodes on the Tatau Platform will be able to use consumer GPUs to fit models that are 10 times larger than what current CPUs can process. Our testing shows that for a fraction of the cost, Tatau will achieve the same results in terms of memory usage as a traditional centralized compute provider.

[5] https://energy.gov/eere/buildings/data-centers-and-servers

Technical Overview

The Tatau platform connects computational resources with computational requests. It provides security of computation, information, and payment by using the blockchain and cryptographic commitment schemes.

The Escrow Contract is built on Ethereum to leverage a pre-vetted code-base of smart contracts, and the Tatau Platform uses IPFS (for data storage and transfer) and BigchainDB (to maintain audit records and transmit administrative metadata). The

platform is able to connect to and harness idle GPUs from 3D rendering firms, as well as other high-performance individual setups, like consumer gaming computers.

Platform Components

Supplier Node

The Supplier Node is responsible for taking in a Computational Request and training the data on the model.

A computational node has:

- An identity that corresponds to specific attributes (libraries, operating system, capacity, and computational capabilities). This information is stored off-chain.
- An address that corresponds to its identity as well as its Tatau Token Address for payment.

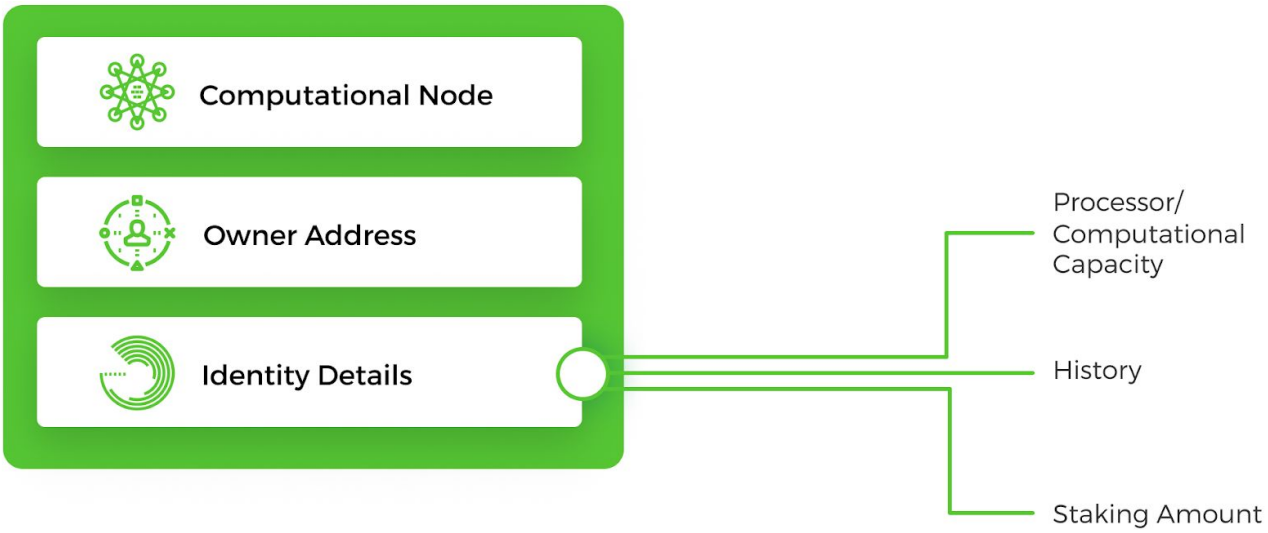

Computational Request

A Computational Request is used to carry the information about the computation that is being requested, as well as metadata such as the software dependencies for the computation and payment parameters relevant to the requested computation.

A computational request contains several key components:

- A public address that corresponds to the Requester. The public address is used for payment and tracking Computational Requests in the network.
- A model, defining the set of algorithms and operations that each Supplier Node will perform on its specific set of data.
- A manifest of software requirements for the model, with corresponding dependencies and languages used.
- A Dataset for the model to be trained on. This dataset will be split into several pieces and stored on IPFS.

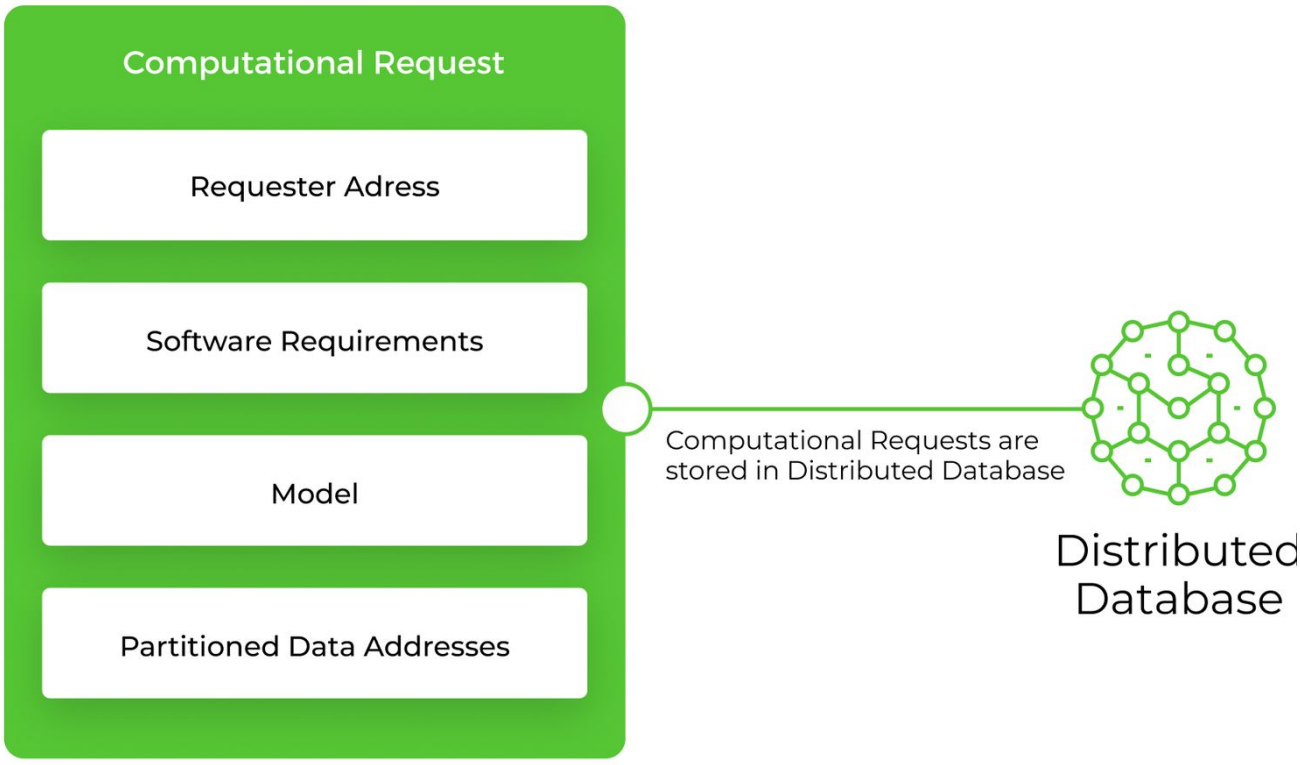

Partitioned Computational Request

For better performance, every Computational Request is partitioned so that the work is performed by many Supplier Nodes in parallel.

Additionally, each Supplier Node receives only a portion of the Dataset, so that no single node has access to the entire Dataset. The data assigned to a given Supplier Node are also encrypted, to prevent unauthorized parties from accessing them.

Escrow Contract

The Escrow Contract connects the Requester and the Supplier Nodes for each Computational Request, and it facilitates payment to the Supplier Nodes.

The Escrow contract contains:

- The unique identifier for the Computational Request.
- The payment offer that the Requester specifies.
- The Address of the Requester.
- Addresses of the Partitioned Computational Requests as well as their associated encrypted computational fingerprints, the payment for each request, the Address of the Supplier Node that has accepted the offer, and the IPFS address of the solution.

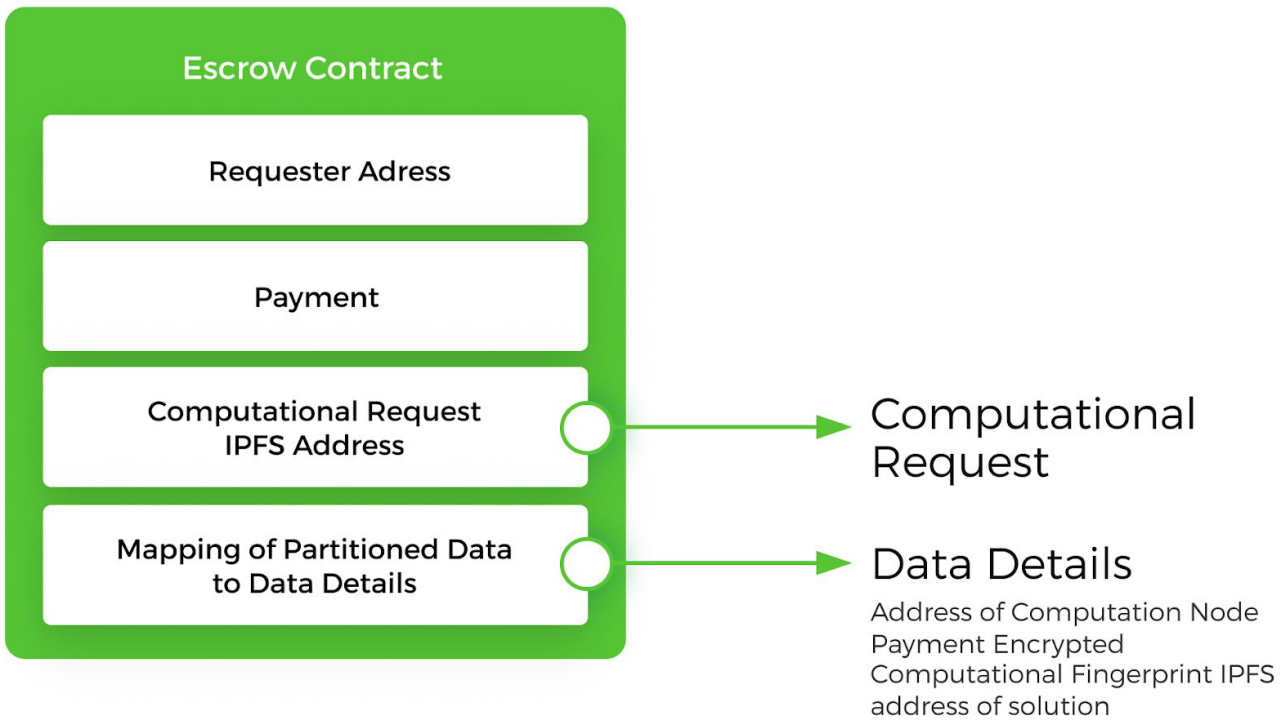

Validators

Validators are special nodes run by Tatau to verify the correctness of the solutions submitted by Computation Nodes. Validators use a proprietary algorithm to check that each result is correct, without having to redo the computation.

Once a Computational Request is completed, the Escrow Contract will only pay out tokens to the Supplier Node addresses once the Validators have successfully verified the results.

Tatau API

The Tatau API provides an interface for interacting with the other components of the ecosystem. The API makes it easy for everyone to access and use the Tatau Platform. The functions of the API include:

- Distribution of data to BigchainDB and IPFS.
- Partitioning and encrypting the Datasets.
- Providing an estimated cost for each Compute Request.

Note that the functionality and features in the Tatau API will expand and contract to fit the needs of the Tatau Platform.

Verification of Computations

Computational Correctness for Video Rendering

Computational correctness for video rendering is verified by checking solutions per each Supplier Node for its set of rendering equations. Requesters can start verification checks that occur outside the Escrow contact. Validators handle and resolve any issue regarding computational correctness.

Computational Correctness for AI Models

Guaranteeing the correctness of an AI computation requires that it be run by at least two Supplier Nodes. To avoid computing each operation multiple times, the Tatau Platform utilizes a Computational Fingerprint that works using probabilistic methods to incentivize Supplier Nodes to output the correct computations.

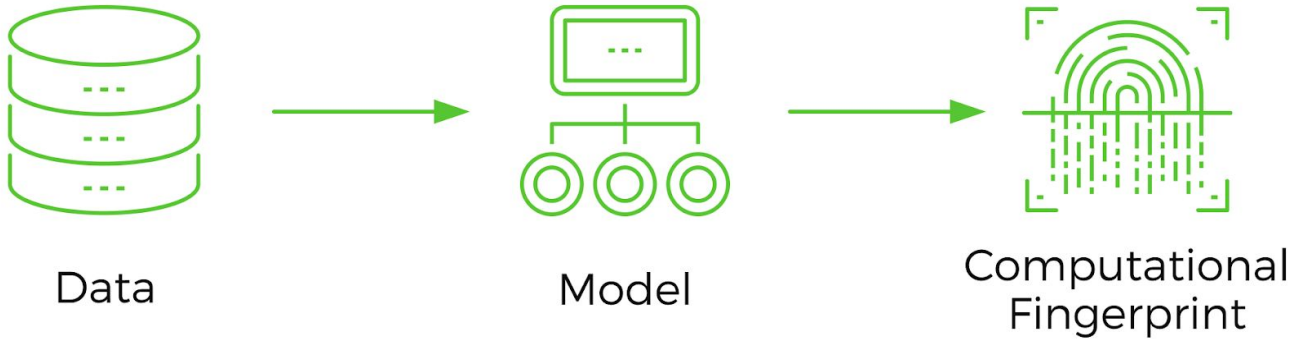

Each partition of the dataset has an associated Computational Fingerprint unique to the stage of the model being run by a Supplier Node. The Computational Fingerprint is evolving as research that the Tatau team has conducted is integrated, but the current implementation will utilize randomness to determine the proper answer.

The Requester will pre-compute the solutions to a subset of data on the model, and hash these solutions together for each partition.

Probability of Catching Incorrect Computation

$$1 - (1 - p)^n$$

Where,
n = number of random samples checked
p = percent of samples that were calculated incorrectly

Once the Supplier Node generates its solutions, it must then calculate the computational fingerprint. This is done by running a piece of pre-compiled binary present in the Model. This pre-compiled code will generate a hash from a random set of the solution.

Using this method of computational fingerprint, it is possible to have a high probability of catching incorrect computations by only running a smaller subset of the data.

Computational Request Life Cycle

Distribute Request to the Network

Using the Tatau API, the Requester can split up their data onto IPFS Nodes as well as create a Computational Fingerprint as a part of their model. Once the model is uploaded to IPFS, the Requester can set up the Escrow Contract with the offer and other relevant information. Once the Escrow Contract is on the Tatau ecosystem, Supplier Nodes can accept the contract.

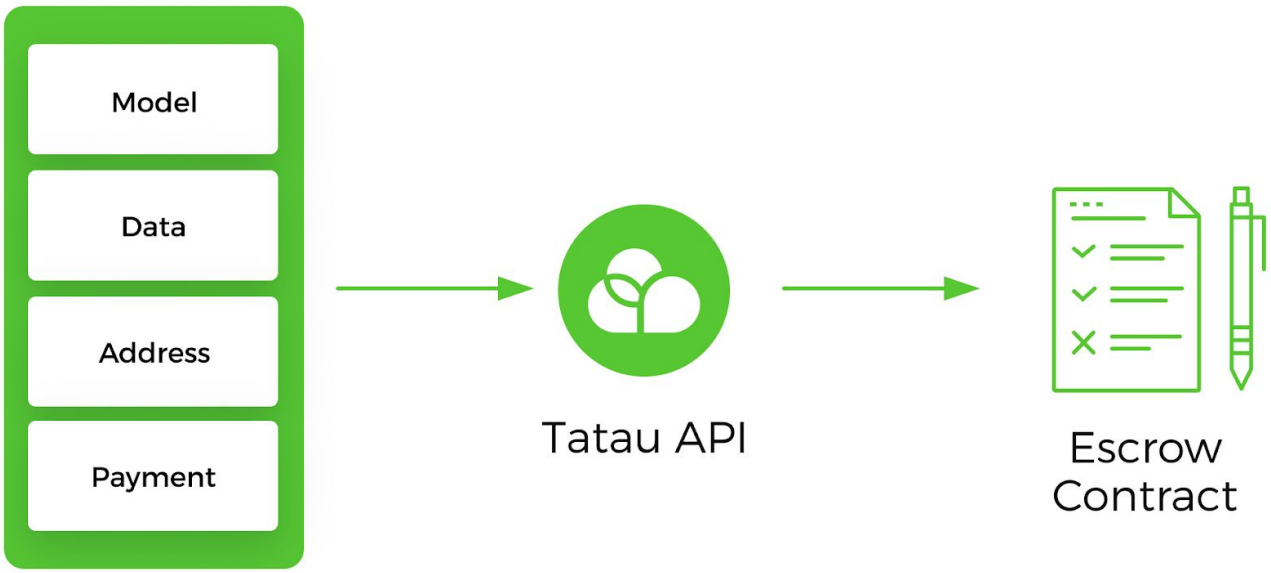

Supplier Nodes Compute Request

After accepting the contract, the Supplier Nodes will then run the model on their set of data, saving this data in another storage node that does not contain the data they are using.

Bring Results Together

Once they have the results of the computations, the Supplier Nodes will share the hash of the data with the Contract, which notifies it that they are done with this batch of computations.

Verify Correctness of Computation

After the Computation Request has been fulfilled, the Supplier Node can submit the computational fingerprint to the Escrow contract, which can verify that this is indeed correct by checking that the Computational Fingerprint corresponds to the private key

for the public key of the Computational Request. Then the Escrow contract can pay the Supplier Node. If this is correct, then the Supplier Nodes will update the model, otherwise a dispute will be handled by the validators.

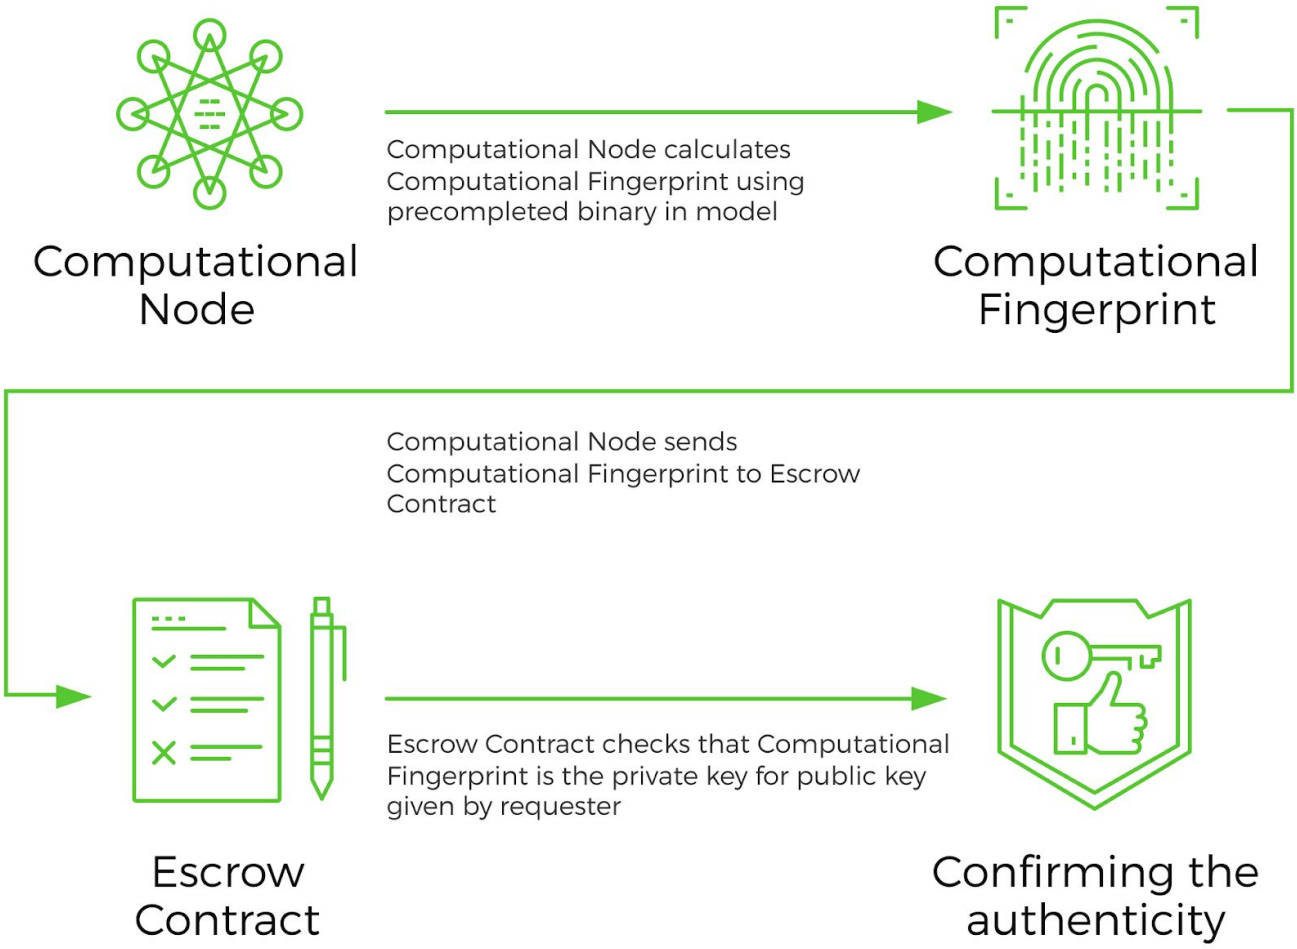

Update Model

After the verification step has been completed, the Supplier Nodes can update the model using their computations. The update can be done in any order since it will include all the data. The update is done when all of the Supplier Nodes have edited the model contract. This can be tracked by the Escrow contract or by the model itself, with the Requester initiating the call to continue to the next step. This verification process can be automated for the Requester, but initially it is designed so the requester can control the entire operation.

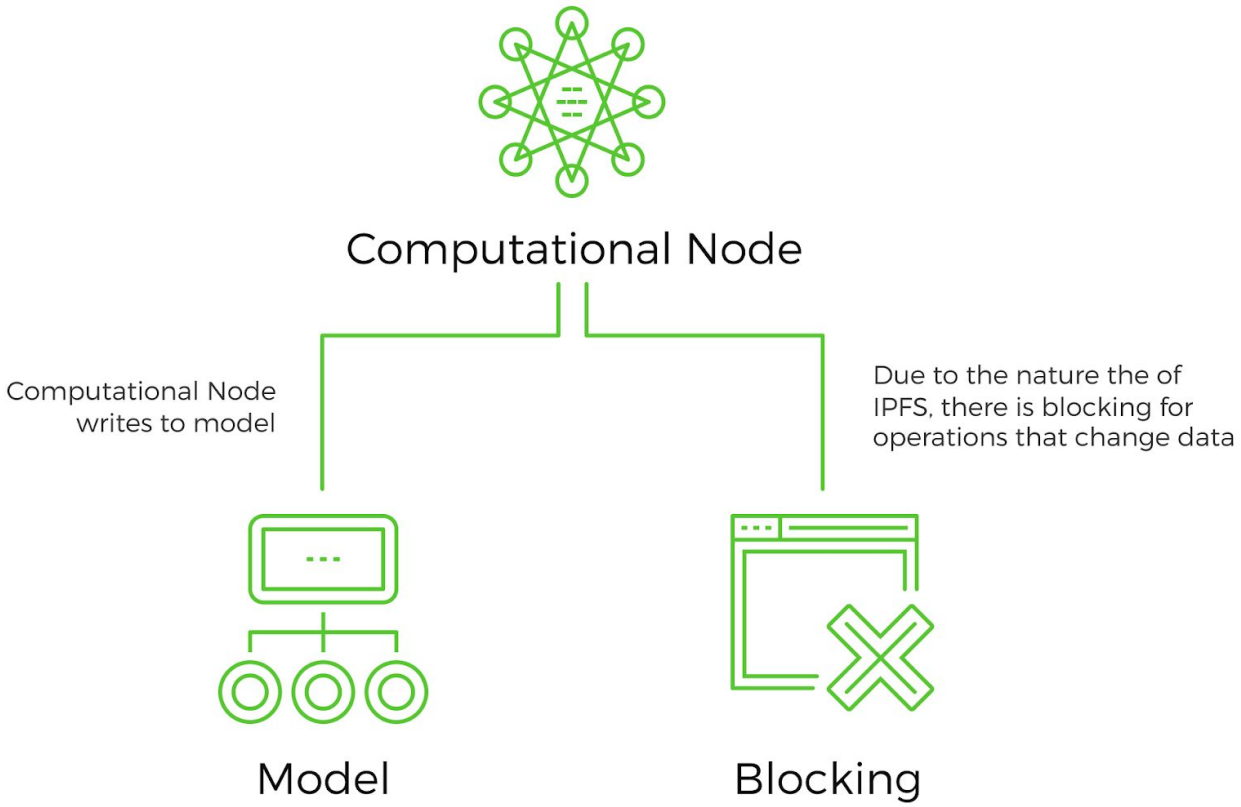

Closing out the Escrow Contract

The process of iterating and updating the model is done until the model is fit according to its specifications, or the Escrow Contract runs out of funds to pay the Supplier Nodes. If the Escrow Contract runs out of funds, the Computation Request is paused until the Requester deposits additional funds into the contract. Once the request has been completed, the Escrow Contract refunds the Requester any unused funds and closes out the computational request.

Handling Disputes

In the case of a potential issue around the correctness of computation, the Supplier Node will be removed from the work queue, and a group of Validators will be called to settle the dispute. The Validators will check to see if the computational fingerprints match and determine whether the Requester made any errors in computation. Once the process is done, the Validators will take the staking that the Supplier Node put into the Escrow contract and remove the Supplier Node from the contract. Further repercussions may apply if necessary.

3rd Party Dependencies

Interplanetary File System (IPFS)

The distributed storage capabilities of the Interplanetary File System (IPFS) will be utilized as a method for lead storage and tracking. IPFS uses the Base58 hash of a file as an address for where it exists. It is one of the most promising, efficient, and reliable decentralized storage protocols in the space. Moreover, IPFS allows for high throughput for data because the data can be split among many different nodes in the network and streamed from all of them simultaneously.

BigchainDB

In addition to large sequences of binary data, the Tatau platform will need to store and query structured metadata, as well as a reliable audit log for system operations. A consortium BigchainDB chain works particularly well for this use case. Its blockchain-based architecture ensures that all updates are tracked automatically, and that updated data propagate efficiently. BigchainDB uses Tendermint for consensus, providing the requisite security without compromising performance.

Proof of Authority Network (POA Network)

The TATAU token is deployed as an ERC-20 compliant contract on the Ethereum mainnet. Any contract that manages large amounts of wealth must be deployed to a platform that provides a high degree of decentralization and security, which Ethereum certainly does.

However, the scalability trilemma[6] indicates that these qualities can only be achieved at the cost of performance. In the case of Ethereum, this means an average throughput of about 15 transactions per second. For a standard ERC-20 contract, this is fine, but the Escrow Contract has significantly higher performance requirements.

In selecting an alternative platform for the Escrow Contract, we were still beholden to the scalability trilemma, meaning that we needed to sacrifice either decentralization or security in order to achieve the necessary throughput. Reducing security was a non-starter, but switching to a somewhat more centralized solution – that is, a permissioned ledger – was palatable.

We opted to use Proof of Authority Network. Miners on the POA Network will initially be owned by Tatau, but over time additional trustworthy entities will be added. For safety, any token transfers that the Escrow Contract performs on the POA Network remain isolated from the Ethereum mainnet; a set of bridge contracts are required to effect these

[6] https://blog.bigchaindb.com/the-dcs-triangle-5ce0e9e0f1dc

transfers, which Tatau can pause in order to mitigate damage in the event that the POA Network is compromised.

Implementation Considerations

Storage of Data

Since the quantity of data will be very large in some of the computational requests, it may be necessary to rethink the data transfer process. One potential solution would be to have the Requester split the data onto multiple occurrences of IPFS to allow Supplier Nodes to load the information in parallel. While a centralized storage system may have beneficial impacts, IPFS was chosen because the same data can be loaded much more quickly once it has been added to the network.

Correctness of Computation

Even though Tatau provides multiple obstacles and disincentives to dissuade people from doing computations incorrectly, it is still possible for a malicious actor to add fake observations into the data and spoil the resulting model. If incorrect computations become a significant problem on the Tatau Platform, it is possible to increase the number of samples that are computed multiple times thereby increasing the likelihood that a malicious actor would be caught. It is important to note that it is strictly impossible to ensure that a computation is done correctly without having to compute the same results multiple times.

Privacy of Data

Whenever big data is involved, one potential issue is ensuring that the data being analyzed remains private. Issues with data privacy can be remediated by cleaning the data (the normal operation of converting categorical to numerical and normalizing data).

Validator Nodes

There may be some concern that the Validator Nodes will act maliciously, which is why they have been given associated identities, are required to stake, and disputes are validated in groups. These precautions make it difficult to be malicious, and malicious actors can be identified and banned as network validators.

Supplier Nodes

There are a few implementation considerations regarding the Supplier Nodes:

Issues Regarding the Computational Fingerprint

There could be an issue where the Supplier Nodes decode the pre-compiled binaries that perform the check of computation. While this could happen, it is difficult to decompile these binaries, and if it does happen, another verification check can be added after the computation is done. This verification would involve testing a random number of samples on the model, and ensuring that they are the same as the ones provided by the computational node.

Timely Computational Request Completion

There is the potential that Supplier Nodes would drop requests or not complete them in the correct amount of time. This could be solved by ensuring that a processor or computational capacity is tracked as a part of each Supplier Node's identity.

Differing Operating Systems and Libraries

There must be a way for nodes to remain synced with current distributions, libraries, and code bases, and they must share the same operating systems. Moreover, since some software (most notably for 3D rendering) is not open source, it may be necessary to qualify Supplier Nodes by the software and packages that they support.

Since both of these implementation considerations rely on keeping information as a type of identity, it may be necessary to track identity and reputation as a part of the Tatau implementation. This could be implemented in off-chain programs or featured in the Tatau dashboard.

Competitive Landscape

Artificial Intelligence Computation

The “verification challenge” for distributed AI computation

The challenge with distributed computational platforms, especially when computing non-deterministic tasks like artificial intelligence, is verification of the fact that (a) the computing has been undertaken, and that (b) it has been done correctly. This verification is required so that the requester (buyer) knows that it is paying only for computations that have been validly undertaken as part of the job, thus triggering payment by the requestor and satisfaction that the job has been properly completed.

The Tatau solution

Tatau is the first platform that will natively undertake “proof-of-computation” for non-deterministic tasks like artificial intelligence processing. Tatau achieves this by using code fingerprints that sit within the job output and can only be verified by the requestor and the Tatau platform.

In competitors’ solutions, the verification is part of the requester task and can not be verified for non-deterministic tasks. To verify, the task would need to be repeated, requiring at least 2x more resources than with the Tatau platform. This means that Tatau is the only distributed platform that enables distributed AI processing to be cost-effective and commercially viable. Other solutions cost more than when using a simple cloud solution like Amazon for the same non-deterministic task.

Rendering

Accessing distributed rendering with commercial-grade software

When undertaking rendering, off-the-shelf commercial software is required to be licensed and installed on each node that undertakes the rendering tasks. Different rendering tasks require different rendering software. Thus different requesters will require access to distributed processing using distributed hardware with different rendering software.

Tatau’s platform will manage the distributed resources and nodes and match the requesters’ jobs with only those distributed nodes that are running the requisite rendering software for that job.

Tatau will connect other rendering farms to the Tatau platform to enable distributed rendering using the required commercial software.

Tatau's competitors only provide for open-source software and proprietary software, which is not suitable for large-scale commercial rendering. Only Tatau provides access to distributed rendering using commercial software.

Operating System Agnosticism for GPU Access

Many distributed GPU nodes, whether in large-scale farms or individual sites, use a version of Windows for their operating systems. Tatau's competitors only support Linux-based systems and cannot leverage the large-scale Windows deployments.

Tatau is OS agnostic, giving it the ability to harness a larger market of available distributed computational resources and therefore Tatau will relate to a much larger user base, which was previously disenfranchised by the early market entrants.

Harnessing Distributed Platforms at Scale

The challenge

The challenge is how to practically manage and harness into the platform hundreds of thousands of unique, distribute, heterogeneous nodes with different operating systems, different GPUs, different software scattered around the world.

On other platforms the requester (buyer) needs to find, set up, and run separate "instances" themselves, having to take care of parallel computing, etc. This is rather complicated, laborious, and can realistically only be achieved by a very small portion of potential clients.

The Tatau solution

Tatau will provide software-as-a-service to support the requester (buyer) with transparent access to distributed resources. The requesters (buyers) therefore do not need to manage each machine. They will run the tasks in the Tatau platform and get the results, without additional management steps related to managing variable and multiple instances and nodes as part of setting up the computation/rendering task.

Next Steps

Once the Tatau Platform has been initially implemented. There are several future steps to make Tatau more versatile and easier to use.

- Expanding distributed computing to include AI, Machine Learning, and 3D rendering, as well as offering different parallelization techniques for different computations.
- Integrating a connection to central cloud storage and optimizing the data storage and transfer mechanisms of the current network.
- Allowing for public data to be shared and queried in the Tatau network to allow for higher-throughput data transfer with IPFS.

The Tatau Token (TATAU)

The Tatau Token is used to in two ways in the Tatau ecosystem:

1. Payment
   The Tatau Token is used when buying and selling computational output on the network. It will also be used to pay fees to parties in the network (i.e. Validators).

2. Staking
   The Tatau Token is used to stake Supplier Nodes and Validators in the network. When mixed with a slashing protocol, this staking method strongly incentivizes honest participation in the network.

Token Generation

Token Sale

Role of Token	Facilitate data processing on Tatau Platform
Symbol	TATAU
Supply	7,500,000,000 TATAU Tokens
For Sale in Pre-Sale	Up to 1,111,111,111 TATAU Tokens (14.8% of minted TATAU Tokens)
For Sale in Public Sale	Up to 222,222,222 TATAU tokens (2.96%, included in Treasury)
Emission Rate	No new tokens will be created
Sale Period	September-October 2018
Accepted Currencies	USD, ETH, BTC
Maximum Goal	USD 10,000,000

Token Distribution

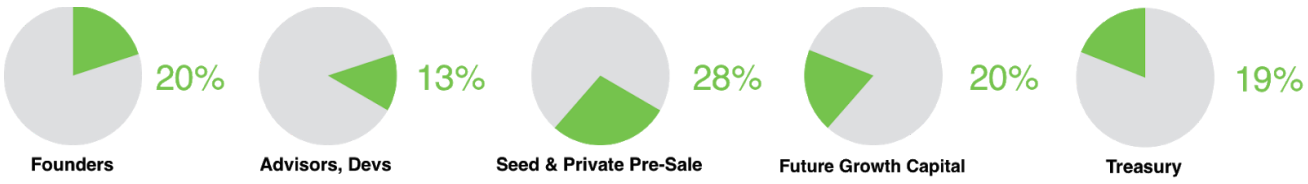

- 1,111,111,111 TATAU tokens (14.8% of minted TATAU Tokens) will be offered for purchase as part of a private pre-sale under the ticker symbol TATAU. 1,000,000,000 TATAU Tokens (13.3% of minted TATAU tokens) have been offered to selected purchasers as the seed round for Tatau.
- 1,500,000,000 TATAU tokens (20% of minted TATAU Tokens) will be allocated to Tatau founders with a one year lock-out period and then a 24 month vesting period. 500,000,000 TATAU Tokens (6.7% of minted TATAU Tokens) will be allocated to advisors and partners with a further 500,000,000 TATAU Tokens (6.7% of minted TATAU Tokens) allocated to developers (including future staff) with sales restrictions to be put in place by management.

- A minimum of 2,666,666,667 TATAU Tokens (35.6% of minted TATAU Tokens) will be held in treasury/reserve by Tatau for clientele growth, business development, academic research, and future raise and market expansion. This will be split into 1,250,000,000 TATAU Tokens (16.7% of minted TATAU tokens) that will be used for Future Growth Capital, and 1,416,666,667 TATAU Tokens (18.9% of minted TATAU tokens) to be held in Treasury with a 24-month holding period. This amount includes the up to 2.7% of TATAU tokens that is are earmarked for a public sale.

Fund Allocation

Funds raised will be allocated as follows. To mitigate volatile cryptocurrency markets and to provide Tatau Limited with a runway of 24+ months, raised funds will be hedged across multiple currencies.

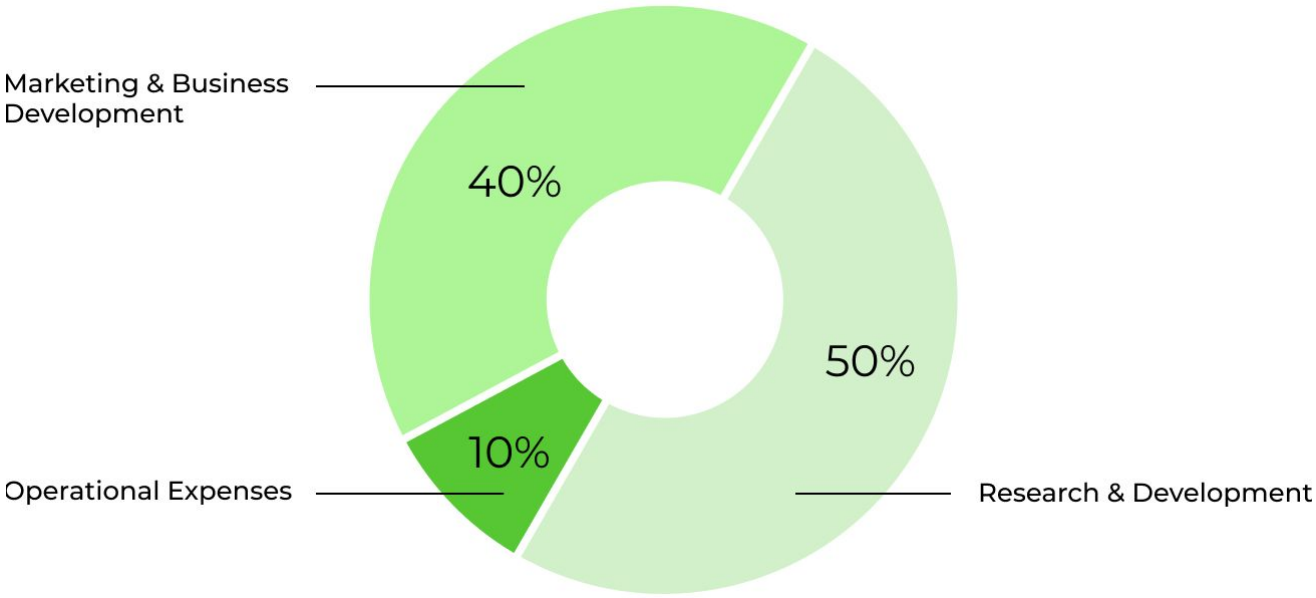

Growth Plan

The Tatau team plans to drive the expansion of the Tatau platform by driving demand to the network, and ensuring that there is enough supply to prevent any shortages or price gouging. By incrementally balancing supply and demand in the network, both customer and buyer can be satisfied. There is no shortage of work, and the work can be done in a timely manner.

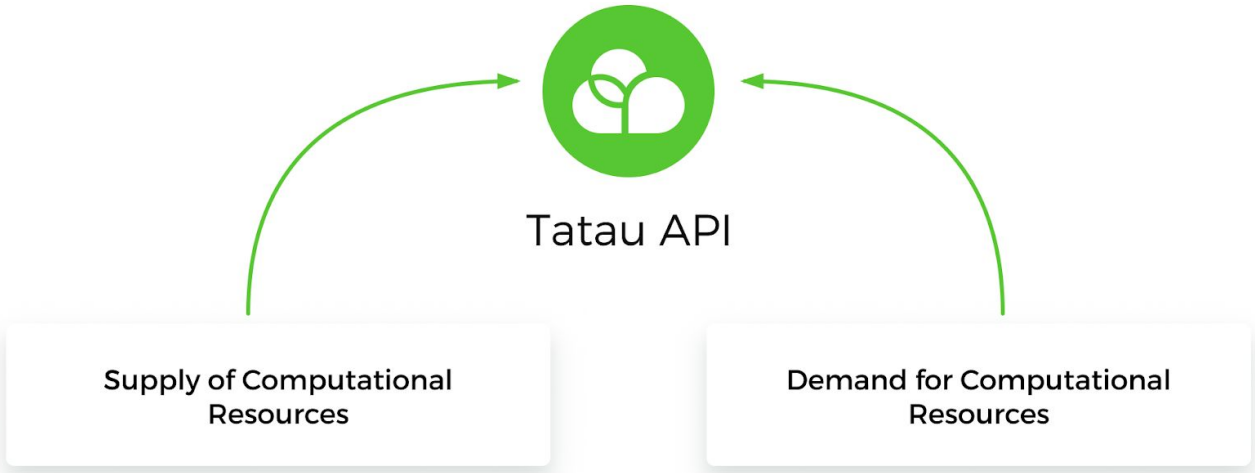

Supply Creation

Onboarding of GPU Supplier Nodes

In order to get a level of initial supply of computational, specific efforts will be placed into marketing the platform to onboard new Supplier Nodes.

Existing supply channels

We see the Tatau platform as a value option to larger cryptocurrency miners, looking to diversify from mining operations, as well as to the significant long-tail of personal miners.

We believe there is both a financial impetus to switch focus, as well as the fact that processing through the Tatau platform is a value-accretive activity.

We will also be working closely with other large-scale existing players such as existing rendering farms, under-capacity GPU data centers, etc.

New supply

We will be stimulating new supply at both the consumer and business level.

From a consumer perspective, we will be undertaking significant social media marketing (both prior to and following the Token Generation Event), Crypto Airdrops, and direct communication with XBox/PlayStation owners via a direct log-in.

Business supply is also critical, with more of a direct-relationship approach via incentives, referrals and potentially a business development network which will cover larger existing players and new business supply.

Demand Creation

For the Tatau network to be successful, there must be sufficient demand for the supply that will be introduced to the system. Demand will initially come from a series of pre-agreements that will be made with 3D Rendering companies as well as companies with significant AI computational requirements.

3D Rendering

The Tatau platform is ideal for companies creating 3D video rendering for movies like Avatar and Lord of the Rings. The capacity of the Tatau Platform would be an inexpensive and eco-friendly method to reduce the bottleneck of 3D rendering. And, the Tatau Platform would make video rendering more accessible to a wider audience of users.

Artificial Intelligence: Communication companies

These companies have considered building data processing services for themselves, but the process of building in-house systems is costly and time consuming. It is easier to utilize the accelerated processing systems of the Tatau Platform. Initial talks with telecommunications companies have resulted in a desire to utilize the Tatau system for communication analytics. A big part of the growth will be to onboard these companies as quickly as possible to generate immediate revenue and publicity for the Tatau Platform.

An overview of our Go-to-market Strategy is set out below, covering both demand and supply creation strategies, and our broader marketing tenets.

Tatau Go-to-market Strategy

Demand side

- Strategic partnership / direct approach
- Movie/Video game companies
- Other companies across broader industries, including:
    - Fintech
    - Medical (imaging, drug discovery etc.)
    - Defence
- Airdrops/indirect market stimulation

Supply side

- Direct through existing channels
    - Value option to large miners
    - Significant long tail of personal miners
    - Large -scale existing players e.g. rendering farms, GPU data centers, etc.
- Stimulating new supply
    - Crypto Airdrops
    - Social Media Marketing
    - Focus on both business (e.g. spare capacity) and consumers (with PlayStations and/or Xbox) via direct login

Broader marketing approach to build market and brand prior to launch

- 4-pronged marketing strategy for TGE
    - balanced whitepaper
    - clear and transparent view of team
    - over-invest in community presence
    - 'above-the-line' PR exposure as required

- Tatau will participate in project, conferences, blog, and groups to promote awareness to our platform within key communities (e.g. AI (speciality distributed AI), green data centres, rendering, etc.
- We will give tokens to academic research programs (AI labs in leading universities). This will support research, generate good will in the AI community, and ensure consumptive usage of Tatau tokens.

Partnerships

BlockchainLabs

BlockchainLabs is based in New Zealand with team members around the globe. Its Founders, Paul Salisbury, Mark Pascall, and Fran Stanjar are dedicated and recognized thought leaders around practical applications for Blockchain and smart contracts. They bring a wealth of experience in developing token contracts and auditing existing ones to ensure best practices for security and functionality, creating crowdsale contracts that can achieve goals with minimal network impact, and liaising with exchange and wallet developers to ensure compatibility and integration of the token.

https://www.blockchainlabs.nz/

Techemy

A world leader in tokenization with a track record of leading some of the world's largest ICO/TGE's. Techemy is a best-of-breed partner for any tokenized venture. Recognized as a golden standard and a leading hub for anything blockchain. Recognizing that Cryptographic Assets are a 4th super-class, Techemy buys, builds or invests into infrastructure, which scale this new asset class.

http://techemy.co/

Blockchain at Berkeley

Strategic approaches to implement blockchain technologies. Research with cutting-edge blockchain and crypto technologies. Working with cutting edge blockchain based ventures.

https://blockchain.berkeley.edu/

Tatau Software Release Schedule

April 2018 Release 1	• Customer and client side application with AI distribution queue layer on top TensorFlow with GPU support
May 2018 Release 2	• Proof-of-computation pilot adding to customer and client side
August 2018 Release 3	• Smart contracts for AI computation tasks and AI proof of computation • Customer Client software with blockchain tasks queue layers • Tensorflow and distribution support • Validator nodes software for AI • Distributed file system support • Processing transactions with RaydenNetwork and Bridge contracts
February 2019 Release 4	• Launching Tatau on testnet • Begin beta processing commercial AI contracts
March 2019 Release 5	• Tatau API • Computational graph analysis for memory optimizations and 6-10x larger models support
June 2019 Release 6	• Launching AI on mainnet
December 2019 Release 7	• Rendering-as-service. Support for renting commercial rendering software as part TATAU platform • Grid Search support, allowing test of different neural networks parameters in parallel • Torch, Caffe, Theano frameworks support • Adding OpenCL support for rendering and AI tasks • Rendering validators • MXNet support • Auto Machine learning layer (AutoML), provide thousands of neural network configuration and model tests
July 2020 Release 8	• Quantum Tensor network simulations with Qtorch

Team

Co-Founders

Greg Kushnir, Co-Founder

Greg Kushnir is a serial entrepreneur having co-founded eight companies and acted as board member and angel investor on several more. He has 15 years of experience in the internet and mobile space and is an expert in the Blockchain domain. Greg specializes in harnessing disruptive innovation, rapid prototyping and commercialization to turn great ideas into successful businesses. His most recent venture, Lina.Review, is the first and only blockchain platform built specifically for the billion-dollar online ratings and reviews market. Greg also co-founded Smart Links Swiss, a Blockchain ecosystem where developers, businesses, and users can collaborate.

Andrew Fraser, Co-Founder

Andrew Fraser has extensive knowledge of the international business landscape having worked in finance and project management in the United States, Malaysia, Africa, and New Zealand. As Group Manager of Key Projects, and Head of Group Market Strategy for Vector, New Zealand's leading energy infrastructure company, Andrew brings a keen understanding of New Zealand's energy sector and business landscape to the project. Andrew spent eight years at McKinsey & Company. He provides valuable expertise in finance, problem solving, and project management.

Michael Fridshtand, Co-Founder

Michael Fridshtand is a software engineer, blockchain enthusiast, and serial entrepreneur with multiple successful companies under his belt. Michael is the Founder of Codesmart and the CTO and Co-Founder of InstallBrain.

Martyn Levy, Co-Founder

Martyn Levy is an entrepreneur and professional director with experience in building and optimising technology and professional services businesses, capital raising, governance and divestment. Martyn is a Co-Founder of Acurix Networks, a managed telecommunications services provider. Prior to that he was Head of Strategy & Business Development for mobile network operator 2degrees as well as CEO of RoamAD, a WiFi equipment vendor, which he shepherded from start-up through to exit. Martyn has also worked in the venture capital industry, been a management consultant, and worked as a banking & finance lawyer in Moscow and Auckland.

Yochay Kiriaty, Co-Founder

Yochay Kiriaty is a Senior Manager at Amazon. In his role, he leads the Alexa Identity and Presence teams as part of Alexa OS. He was also a Principal Program Manager at Microsoft Azure responsible for growing the service from incubation stages to a multi-hundred-million-dollars business. Yochay is an expert at successfully managing large scale projects, starting from the early stages of the design through the different incubation and development stages, and culminating with the delivery, integration and go to market.

Advisors

Jeff Pulver, Advisor

Jeff Pulver is an American Internet entrepreneur known for his work as founder and chief executive of pulver.com and co-founder of Free World Dialup, Vonage, MoNage, Alchemist, and Zula (app). Pulver has been called a Voice over Internet Protocol pioneer, and has written extensively on VoIP telephony, and the need to develop an alternative to government regulation of its applications layer. As Vice Chairman at Alchemist, his current focus is the intersection of computing, AI and communications, where his team have advised organisations such as Ethereum, tZERO, RSK and Factom. Pulver has also founded the Security Token Association, a new trade association whose purpose is to contribute to the global growth of the Security Token industry.

Dr Michelle Dickinson (MNZM), Advisor

Dr Michelle Dickinson MNZM, also known as Nanogirl, is a nanotechnologist and science educator. She set up and runs New Zealand's sole nanomechanical testing lab, and is a senior lecturer in Chemical and Material Engineering at Auckland University, and an associate investigator at the MacDiarmid Institute for Advanced Materials and Nanotechnology. She was a recipient of the Prime Minister's Science Communication Award and Royal Society's Callaghan Medal, was made a Member of the New Zealand Order of Merit (MNZM) for services to science, and was named a Women of Influence for Innovation and Science.

Dr Yaniv Gal, Advisor

Dr Yaniv Gal (PhD) is a Research and Technology leader with over 20 years of experience in both commercial and academic environments. He is currently Director of Artificial Intelligence at Ohmio, a leading autonomous vehicle company. He has extensive experience in building, managing and mentoring research teams in the development of new technologies, in the fields of artificial intelligence, machine learning, computer vision, medical imaging and related multidisciplinary fields. Dr Gal is the author of over 40 peer-reviewed academic papers and inventor of over ten patents, and played a key role in building world-leading technologies, directly increasing the value of companies worked with/for.

Nimrod May, Advisor

Nimrod May is one of the leading marketing professionals in the Blockchain space, with a global reputation for thought leadership and marketing execution. He is currently the Chief Marketing Officer at SIRIN Labs, where he led the marketing function supporting their $157.8M token raise, as well as their ongoing marketing efforts including the Brand Ambassador relationship with Lionel Messi.

George Samman, Advisor

George Samman is a leading Blockchain consultant, investor and educator. He consults to financial institutions globally on Blockchain, including KPMG and Brian Kelly Capital Management, and has led research on how blockchains and distributed consensus systems are applied to financial technologies. George also co-founded BTC.sx, now magnr, a bitcoin trading platform. He writes a blog on blockchain technology and use cases at sammantics.com.

Paul Salisbury, Advisor

Paul Salisbury has been working in the blockchain sphere since its infancy. He is a recognized thought leader around practical use cases for blockchain and smart contracts, and an expert at helping decision makers navigate the blockchain landscape. Salisbury is a Co-Founder of Blockchainlabs.nz.

Giotto De Filippi, Advisor

Giotto De Filippi is an experienced Blockchain practitioner, and one of the industry experts in the field of Tokenomics. He is the Co-Founder of JUR, a revolutionary blockchain-based dispute relationship system. He is also a respected advisor to multiple companies in the Blockchain space, including Decentraland, Steamr, Starbase and Cardstack.

Mitchell Pham, Advisor

Mitchell Pham is the Chairman of the NZ Technology Industry Association (NZTech) and a member of the NZ Financial Innovation & Technology Association (FinTechNZ). Internationally, he is recognized in leadership as a World Class New Zealander by KEA Global, an Asia 21 Fellow by the Asia Society, and a Young Global Leader by the World Economic Forum. Pham is a Co-Founder and Director of the Augen Software Group in New Zealand. He is also a Co-Founder of Smart Links Swiss, a company solving real world problems with blockchain technology, and Lina.Review.

Rich Elliott, Advisor

Rich is an entrepreneur with over 18 years' experience in 3D visualisation, rendering and animation. He has proven experience in building companies from start-up to multi-million-dollar revenues. Recently he built a 2D/3D animation studio specializing in pre-production services for Film & Television. Rich is currently developing several innovative cryptocurrency related projects, centered around GPU-based processing.