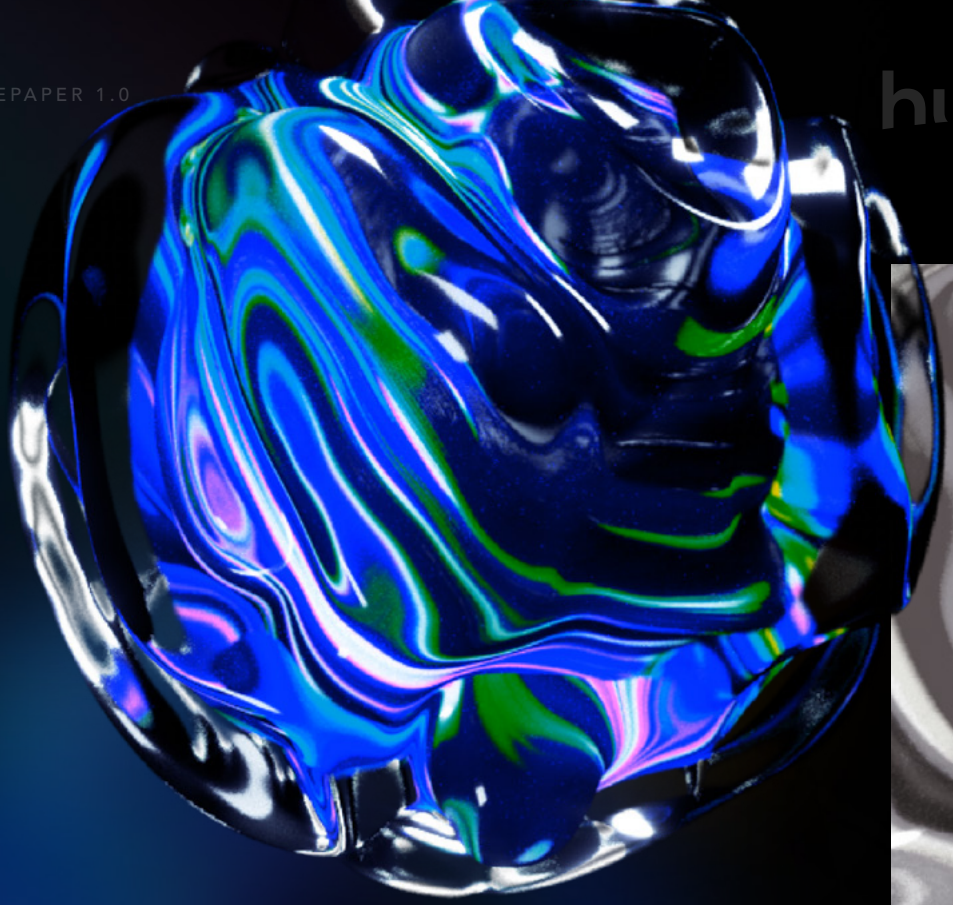

humans

HEART DRIVEN AI

WHITEPAPER

DOCUMENT VERSION 1.0
DECEMBER, 2021

TABLE OF CONTENTS

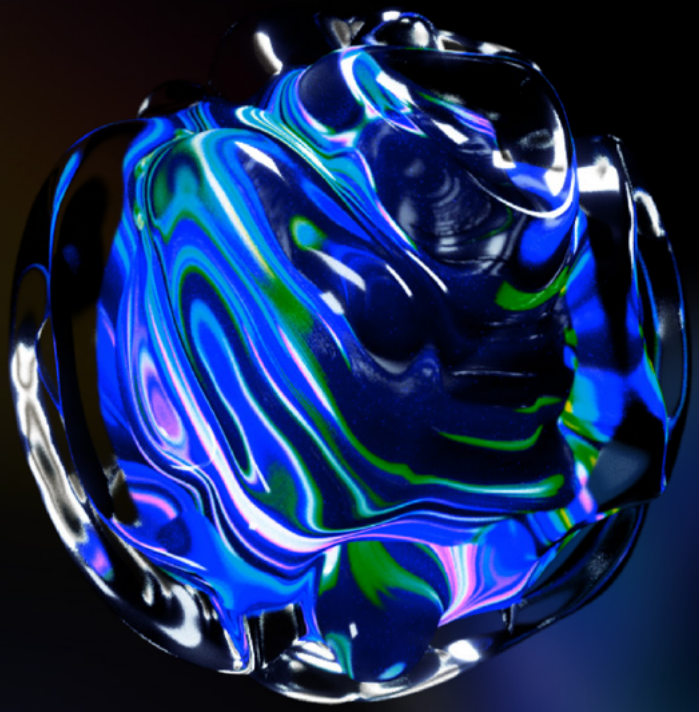

Legal disclaimer:

The Whitepaper is intended for general informational purposes only and does not constitute a prospectus of any sort, or a solicitation for investment or investment advice, nor does it in any way pertain to an offering or a solicitation of an offer to buy any tokens.

Nothing in this Whitepaper should be treated or read as a guarantee or promise of how Humans business or the tokens will develop or of the utility or value of the tokens. This Whitepaper is set out to be a guideline for the exploration by Humans of the presented technology, which could change at its discretion, and the success of which will depend on many factors outside Humans control, including financial, technical, regulatory or other reasons.

Certain statements, estimates and suppositions are forward looking statements that are based on and take into consideration certain known and unknown contingencies which may cause the estimated results to differ factually from the features presented.

The information herein may not be exhaustive and is subject to change.

1

Introduction

1. Introduction

1.1. Abstract

Humans is AI — with heart.

Introducing the first framework for ethical AI and blockchain, Humans creates an all-in-one platform for AI-based creation and governance at scale, beginning with an initial focus on synthetic media.

Through its Humans studio and token-based ownership and accountability system, Humans is designed to ensure contributions are fairly rewarded and that every AI is kept honest over the long term.

Humans is a next-generation blockchain platform that brings together an ecosystem of stakeholders around the use of AI to create at scale.

It combines a library of AI tools into a Humans Studio suite from which users can pick and choose as they bring their ideas to life.

Through this initiative, individuals are empowered to generate and own their digital likenesses to create innumerable digital assets, synthetic media, AI apps, and other digital assets.

In addition, Humans utilizes blockchain technology to generate Non-Fungible Tokens (NFTs) to ensure transparency, provenance, accountability, and long-term governance.

Is a human behind every AI decision.

1.2. Market overview: The potential to solve existing problems.

Welcome to the Era of AI.

We stand at the dawn of a new era, in which the influence of AI on the data that we generate and consume is unparalleled. Our Google search, Netflix and Spotify feeds are customized and displayed based on our personal preferences. With time, the content itself will be personalized.

The advent of this technology holds great promise in terms of how it can help us save time, energy and money. However, it also poses many challenges, some of which have been identified as critical issues that Humans seeks to address with its range of solutions.

Market Fragmentation

Most AI researchers and developers gravitate to this space due to a passion for creating new technologies. In research or just out of pure enthusiasm, they develop neural networks for which they do not have further use. If they want to attract users to their creation, they are forced either to sell it to someone else or to take it to market themselves, both of which are distractions from doing what they love - building AI technology and applications. For now, AI creators most often make their neural networks available as open-source products. This way, they allow for the further development of AI technology, but the authors do not participate in monetizing the final implementations.

Humans offers AI innovators an opportunity to earn livelihoods by bringing their solutions onto a global stage where others can enjoy the value of their technologies over the long term. Through the blockchain infrastructure, the Humans platform will ensure that every contributor is fairly rewarded.

Market Accessibility

Throughout history, we have used tools to save time and energy and enable ourselves to achieve things on a scale that was impossible before.

Today, AI is poised to help us manage data, automate processes and create media on an unprecedented scale. Still, it remains challenging to use and unaffordable for the vast majority of users. While AI democratization is heavily discussed, and many projects advocate it, most AI benefits are at the sole disposal of large corporations such as Google, Amazon, etc.

Humans offers a user-friendly Humans Studio where anyone can access AIs and use them for an unlimited number of applications, with synthetic media as just the first of many more to come. The community-based approach coupled with blockchain infrastructure will allow users to access the AI products at affordable prices.

Ethical Accountability

Just as with the birth of virtually any other industry, AI enables us to do great things but there is also potential for its misuse for harmful and unethical applications. We have already seen cases of "weaponizing" AI through facial recognition and location tracking, discrimination during automated recruiting processes, internet search results, or data misuse to generate deepfakes, fake news, altered voices or faces. Unfortunately, as of right now, there is little to no accountability for AI's presence in the market and its impact on the world.

The three major areas of ethical concern around AI represent privacy and surveillance, bias and discrimination, and the role of human judgment. Several specialized organizations, such as the Center for the Governance of AI at the University of Oxford, the Human-Centered Artificial Intelligence Institute at Stanford University, and others, research these issues. In addition, the European Commission has appointed a High-level expert group on Artificial intelligence (AI HLEG) to produce guidelines, requirements and recommendations to address them, while adopting a human-centric approach.

Via the human validating system and unforgeable blockchain records, the Humans project has pioneered a solution representing a new ethical era in AI. As AI is utilized to create content and a range of digital assets, we ensure a clear record of who contributed to, and bears responsibility for, every element over the long term. In other words, Humans provides a way of keeping AI honest.

1.3. AI use cases

Currently, we are at a tipping point regarding the power of AI capabilities. These technological innovations will soon change how we interact with communication, entertainment, education, healthcare, banking, insurance, and other industries.

Synthetic media are the first to enter this new technological era. Synthetic media, or AI-generated media, is an all-encompassing term for the artificial creation or modification of media by "machines" – particularly programs that rely on artificial intelligence and machine learning.

In other words, it is media that is produced by technology. Today, some types of synthetic media include AI-written music, text generation, imagery and video, voice synthesis and more. The field is ever-expanding as synthetic media companies aim to disrupt more and more parts of traditional media, making new things easier to create.

Synthetic media are already a part of our present. Actor Val Kilmer, who lost his vocal cords during cancer, can play again using his synthetic voice. Bruce Willis leases his digital image and voice rights to star in an ad for a Russian telecom.

More and more industries will be joining the AI trend down the road. Following the market, Humans creates the technology, the model of interaction between stakeholders, and the economy around Synthetic media AIs. Next, we will open our ecosystem to Media & Streaming, Marketing & Advertising, e-Commerce, Health, Financial services, Education and other specific AIs.

For example, AI models can be trained and deployed to detect fraudulent financial activity or optimize sales and marketing based on consumer behaviour. Financial services (along with other industries) can integrate trained virtual sales and customer support agents. Online retailers can benefit from AI algorithms to analyze user comments and reviews and customize the online experience for the shoppers, making it even more engaging when transforming plain text pages into short ads showcasing their wishlist products. Marketers can use AI tools to analyze and manipulate voice when interacting with customers, assess user behaviour and sentiment to generate more relevant offers, or personalize the experience of using a product or service. Lufthansa already did it with an AI system suggesting lounge access between flights to travelers whose customer data fits the appropriate profile.

As the next step, with the help of AI, there comes an opportunity to transform data into businesses. Any company that generates data, or is a data provider, can mint their new line of business on the blockchain through a domain-specific chain (more on that below, in *Section 3. Humans Blockchain*).

For example, health clinics can leverage their anonymized research data to train neural networks and generate treatment schemes. They feed the gathered data to the AI algorithm, and the trained AI can produce treatment plans on request. Clinics may further provide the service of generating such programs on demand for other healthcare practitioners. It can help shorten treatment durations when the trained models are applied to the new patients' data. At the same time, client healthcare practitioners or their patients will not have access to the original data used to train the AI.

In the same way, any business generating data valuable for other businesses has the potential for transforming this data into a new product or service, and therefore, into a new line of business.

1.3.1. Synthetic media use cases

The only limit for the real-life application of synthetic media is our imagination. We will be able to speak in real-time in any language, be the main character in our favourite sitcom and scale our creativity and potential. Examples of this include:

Education

Customization of online courses to be delivered in multiple languages or optimized based on audience-specific needs, real-time delivery of education programs to remote or underserved areas by a "digital twin" of an in-person teacher.

Entertainment

Digital DNA creation, a narration of audio-books by one's favourite actors or deceased relatives, replacing actors from a streaming movie to your digital likeness, etc.

Media

Content optimization by generating audio or video shows that are free from mistakes, stuttering, etc., scaling-up media (e.g., podcasts, videoblogs and shows) by portraying the original voice and face of a presenter in multiple languages.

eCommerce

Transforming text product pages into channel-optimized video clip ads, increasing brand engagement and conversion rates.

Advertising

Creation and editing of advertisements without the need for the physical presence of the actors.

Finance

Provision of personalized, user-friendly banking, insurance, and financial services by a digital human replicating the best human-to-human customer experiences.

Business

Pitches done by anonymous avatars to help democratize startups' access to funding; training videos for businesses' employees.

Sales & Marketing

Automatic personalized videos based on customer engagement, customer onboarding, product walkthroughs, demos, sales campaigns, customer support, etc.

Corporate Communication

Turning slide decks into video presentations.

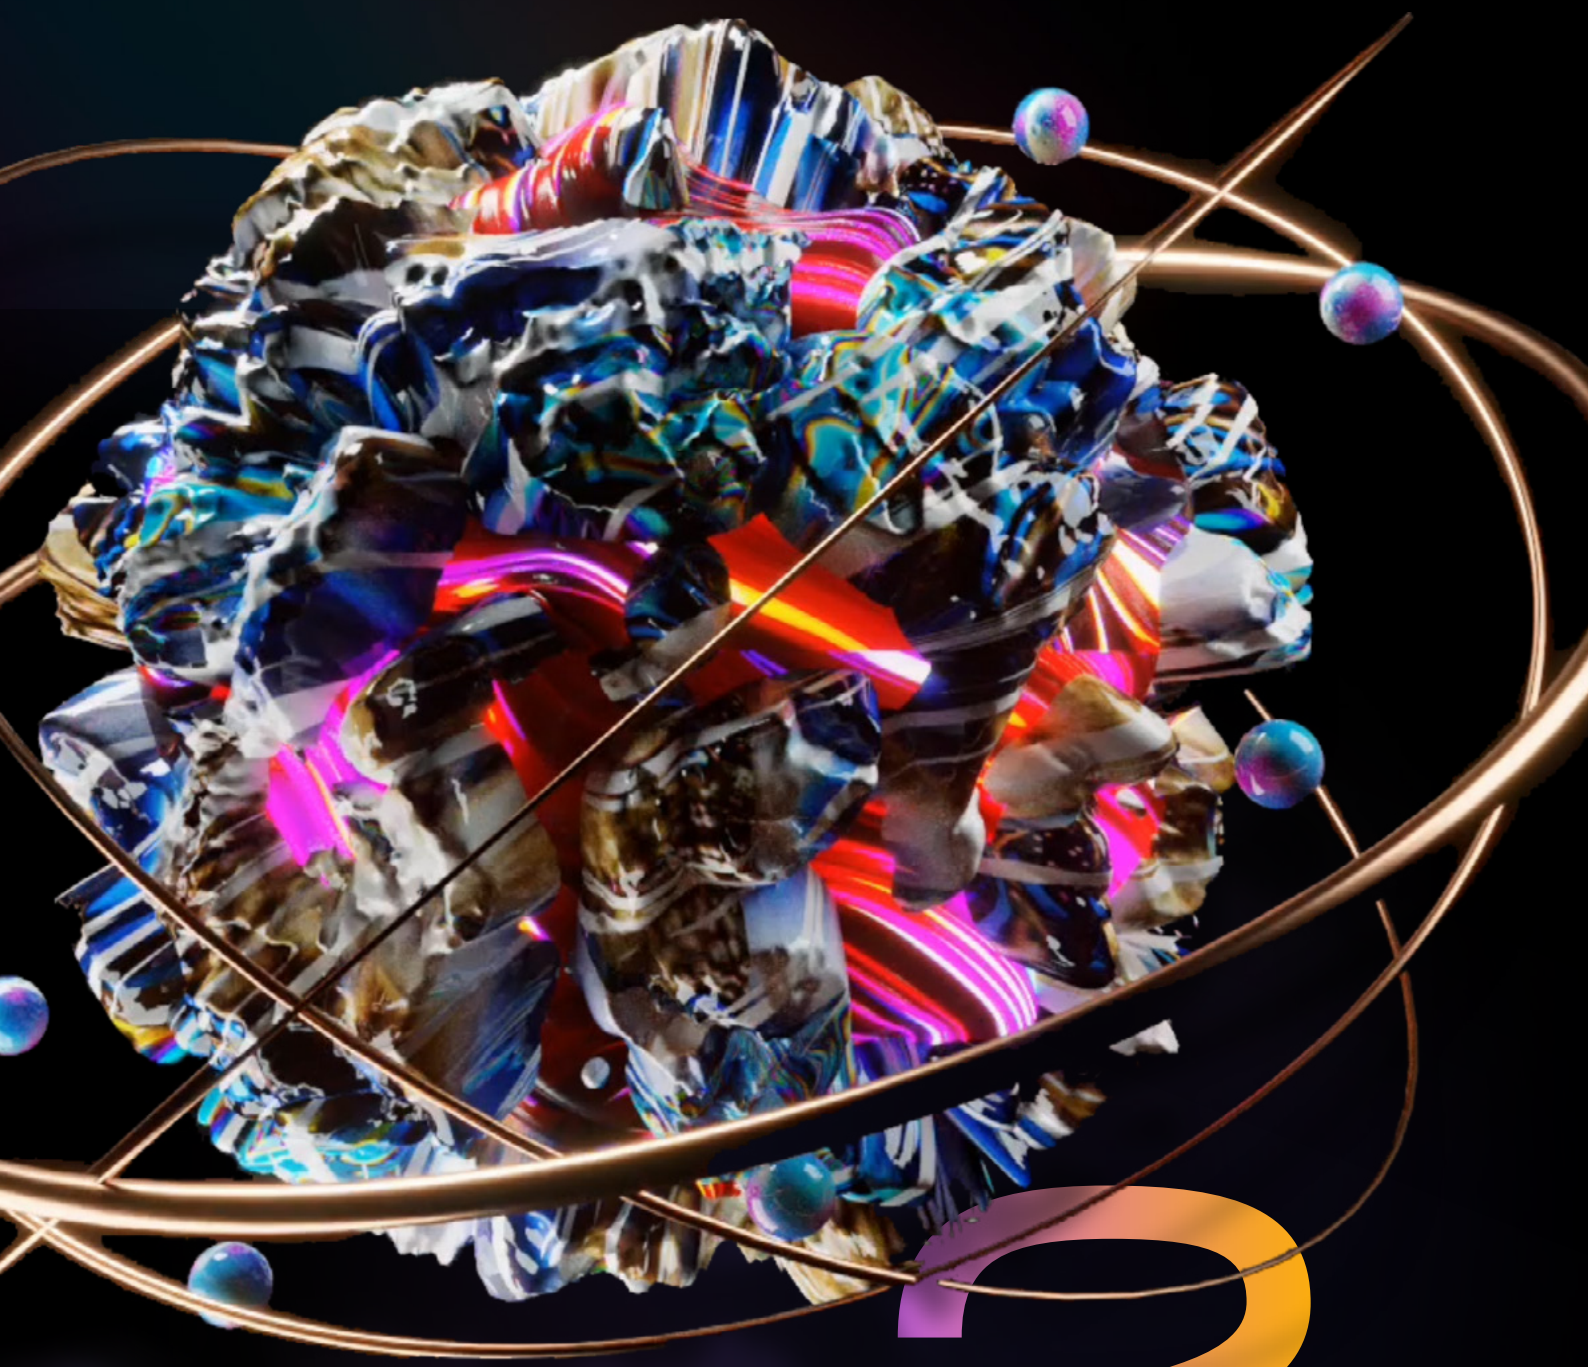

2

Ecosystem

2. Ecosystem

Humans is envisioned as a decentralized, open ecosystem with transparency at its core; an all-in-one technology solution for AI-as-a-Service. It will combine traditional software solutions, blockchain-based record-keeping and governance, and human oversight through the use of biometrics as a biological control key.

2.1. Stakeholders

AI developers

(via universities, research centers, corporations or independent) create AI algorithms and deploy them onto the Humans platform and set conditions for their further use.

Data providers / Users

upload their personal data onto the platform (facial and body images, voice recordings, kinematics videos, etc.) creating their "Digital Genomes" to train different AI algorithms with it. As a basic rule, data providers own all their data. However, they can allow others to use AIs trained with their data and set conditions for their use.

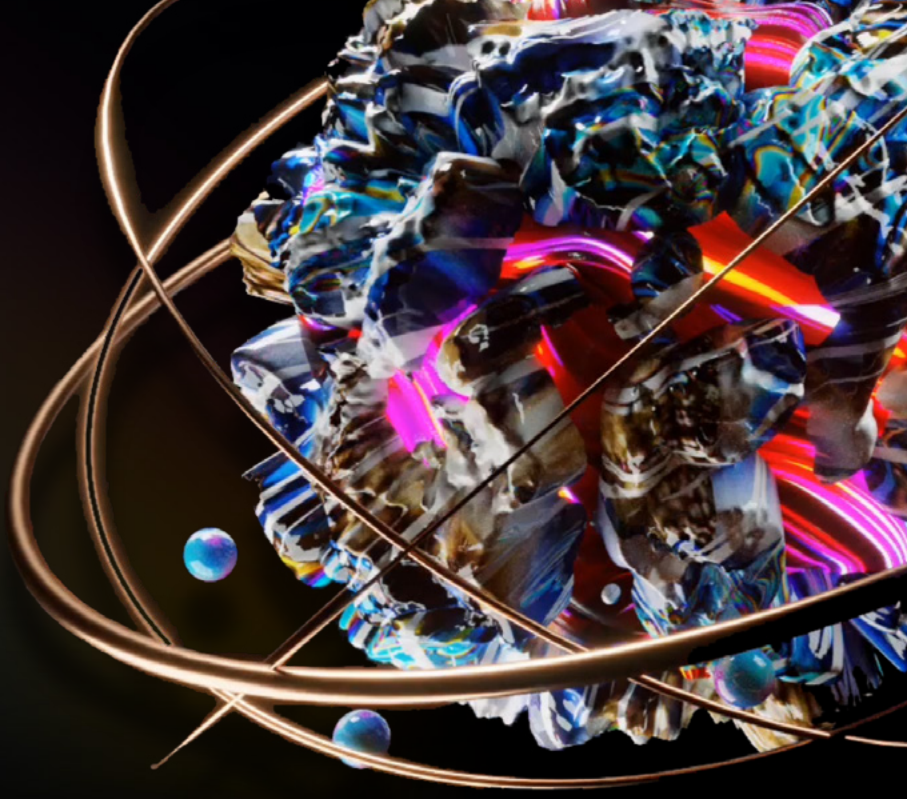

App builders

develop various apps based on AI algorithms provided by AI developers and AIs trained by data providers. They do it by creating workflows of different algorithms and adding more traditional functionality on top of them (e.g., after effects).

Human validators

leverage their biometrics to participate in AI supervision. They act as an independent party, validating all the data uploaded to the Humans platform by data providers and all the requests to use algorithms and content, and ensure that these comply with the requirements set by their owners.

Server validators

secure the network and provide decentralization. They are responsible for validating transactions, and ensuring all parties contributing to a valuable outcome receive appropriate compensation.

2.2. Environments

Humans platform includes specific environments for different stakeholders to work in:

Research Studio
(incl. AI Library)

Humans Studio

App Marketplace

Proof-of-Human
(PoH) App Suite

2.2.1. Research Studio

The Research Studio represents the entry point of AI developers in the Ecosystem. Any AI developer can sign up on the Humans platform, upload their algorithms, train them with their own datasets or datasets available at the platform, test and deploy them, create settings for commercial use by third parties, and publish. Published neural networks become available for the third parties at the AI Library, a "marketplace" of sorts built-in into the Humans platform.

Having access to the Humans community, using their algorithms with their own datasets, the AI developers will get feedback on the AIs' performance (e.g., if lips move precisely enough, if sounds of different languages are pronounced correctly, etc.). Such feedback will allow them to compare their neural networks' performance with the average performance characteristics of similar networks on the Humans platform and further refine their AI algorithms.

After deploying the algorithm, the researchers will set the usage rules for their technology - for example, how the algorithm can be used, who can use it, pricing considerations, and even if governance may be granted to other researchers. Authorship and the rules will be recorded on the blockchain to ensure a fair reward when the algorithm is used. Monetization and access to the global stage will motivate researchers to bring more of their creations to the Humans platform contributing to the further growth of the entire ecosystem.

2.2.1.1. AI Library

The AI Library was built with a set of several standardized ready-to-use AI implementations. Independent AI tech providers can add their algorithms to the library.

The algorithms that are currently available in the AI Library include, without being limited to:

Text to audio

(both in-house developed and cloud-based)

Audio to video

(e.g., speech transformed into a face narrating the audio)

Video quality enhancement

Key points to video, allowing a more granular control

(e.g., each body part is assigned a key point which would be synchronized with the audio)

Real-time face swap or face modification

Speaker identification

(e.g., extracting a piece of audio for a specific speaker)

We are also working on updating the library with metrics that would help assess the quality of the output - for example, the synchronicity and accuracy of voice and lips/overall facial movements, etc.

We aim to make a quantum leap by integrating and developing innovative cross-modal neural models that can seamlessly manipulate and transform different data types. Most of the current approaches are mono-modal, while Humans seamlessly infers from and integrates multi-modal data together with one another (i.e., visual, audio).

As such, we will enable:

Cross-modal processing of textual and visual input to create efficient and reusable representations in a shared space

(e.g., using a few image samples from a user, the algorithm can generate full facial videos with various expressions).

Cross-modal parsing of textual, audio, and visual content for generation of digital media

(e.g., voice cloning, key points to video animation).

Cross-modal generation, like producing visual data from text or audio and vice versa, including mixed content

(e.g., transforming a text from the e-commerce page to an interactive product video).

As a result, only small amounts of input data from a person is required to generate their Digital DNA (Digital Genome). Therefore, it is highly beneficial for the system's scalability and efficiency.

From the methodological point of view, the tentative workflow starts from:

a. Cross-modal data input provided by the user or by inhouse datasets, manual and automatic annotation (data labelling).

b. Use of models for neural/ structural cross-modal representation.

c. Finding new solutions for neural/ structural cross-modal retrieval and generation tasks to benefit from each modality.

d. Evaluating performance in pragmatic use cases.

A special focus will be devoted to fully understanding the capabilities and limits of the models by focusing on their sustainability in terms of efficiency, computational power, the capacity of generalization across domains, transparency, and explainability (achieved by linking latent to explicit representations).

The library will be based on the unifying paradigm aiming to find synergies between supervised neural networks (going beyond current convolutive autoencoders, GANs, Transformer-based neural networks, Capsules and graph-based networks) and symbolic representations, as those obtained from multilingual lexical-semantic knowledge graphs.

The library will enable a trustworthy and reproducible comparison between and the integration of various state of the art methods.

The library will follow the principles of the Nvidia DALI library, supporting a collection of highly optimized building blocks for loading and processing text, images, video and audio data allowing it to be used as a portable drop-in replacement for built-in data loaders and data iterators in popular deep learning frameworks.

Additionally, considering that deep learning frameworks have multiple data pre-processing implementations, resulting in challenges such as portability of training and inference workflows and code maintainability, the data processing pipelines implemented will be portable and could easily be retargeted to TensorFlow, PyTorch, MXNet and PaddlePaddle. It will assure compatibility with some of the existing standards.

The library will support distributed training pipelines with mixed-precision for a computational speedup, faster data loading (e.g., via Nvidia DALI), online linear evaluation for better prototyping, and many other training strategies and tricks presented in recent papers via Pytorch lightning.

AI models may be optimized for specific data inputs. This may occur on an independant/unsupervised basis that requires no labels or a supervised one that does. This training can be automated (~continuous integration-style) or configured manually (data+model), resulting in a trained model taking the form of a binary file.

Users and app builders access the AI Library via the Humans Studio and bring their ideas to life by utilizing any number of AI algorithms and trained AIs. Neural networks for sentiment analysis, text-to-voice, text-to-video, voice-to-video, language translation, data analysis, etc., will all be able to work together in unlimited combinations and be saved as Apps.

2.2.2. Humans Studio

The Humans Studio represents the entry point of data providers/users and app builders in the Ecosystem.

In the Humans Studio, data providers can create their unique Digital Genomes consisting of multi-modal personal data (visual, audial, and other specifics) for use with their creations. As a basic rule, users own all their data. To make data available in the Humans Studio, third-party validators will need to validate the data and authenticity of ownership. Data that has been validated, remains so in perpetuity unless the owner requests changes, which will require another validation.

Furthermore, via the Humans Studio, users can choose an algorithm from the AI Library and train it with their Digital Genomes, as well as play combining different algorithms to create their recipes. The resulting trained algorithms may earn their owners compensation as they are used to produce synthetic media. Requests for use will be submitted for approval by the creator or their designated agents. Everyone who contributed to their creation will get their share of compensation.

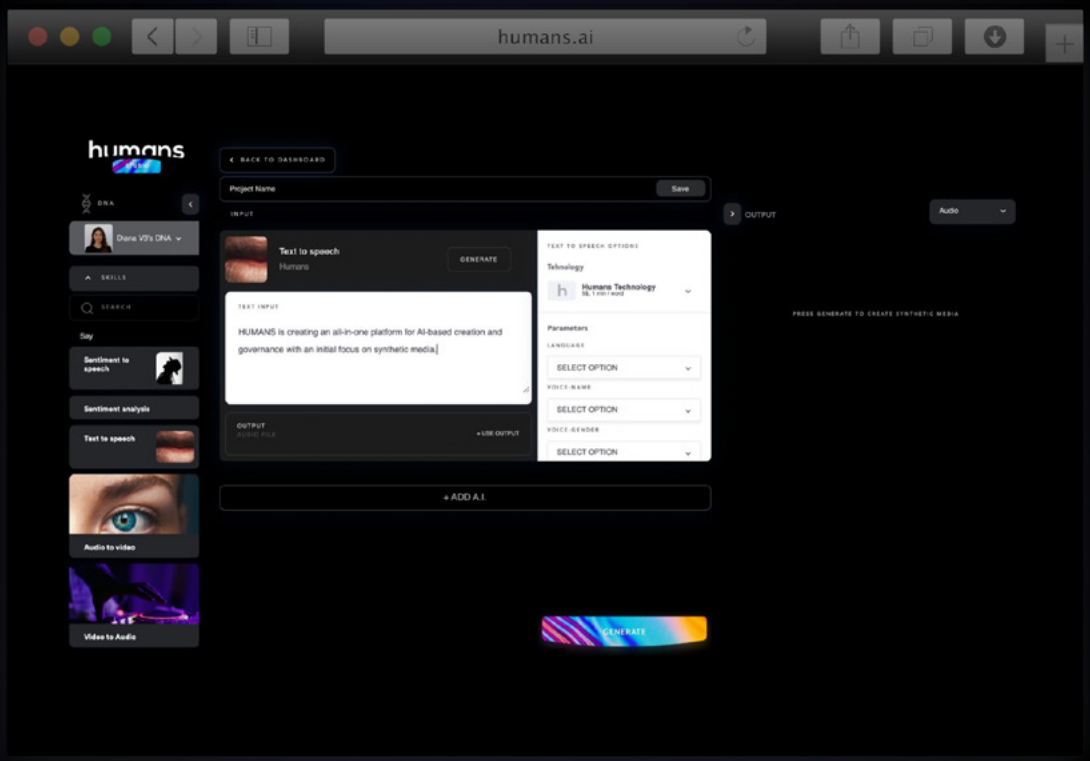

Conversely, users can combine different neural networks for the perfect recipe tailored to a specific use case, as well as browse use cases and corresponding potential AI technologies. For example, users may combine algorithms creating facial movements, voice sounds, and background to create a video of the avatar reading text. Or they may complement such a combination with neural networks producing specific features, such as hyper-realistic lip movements.

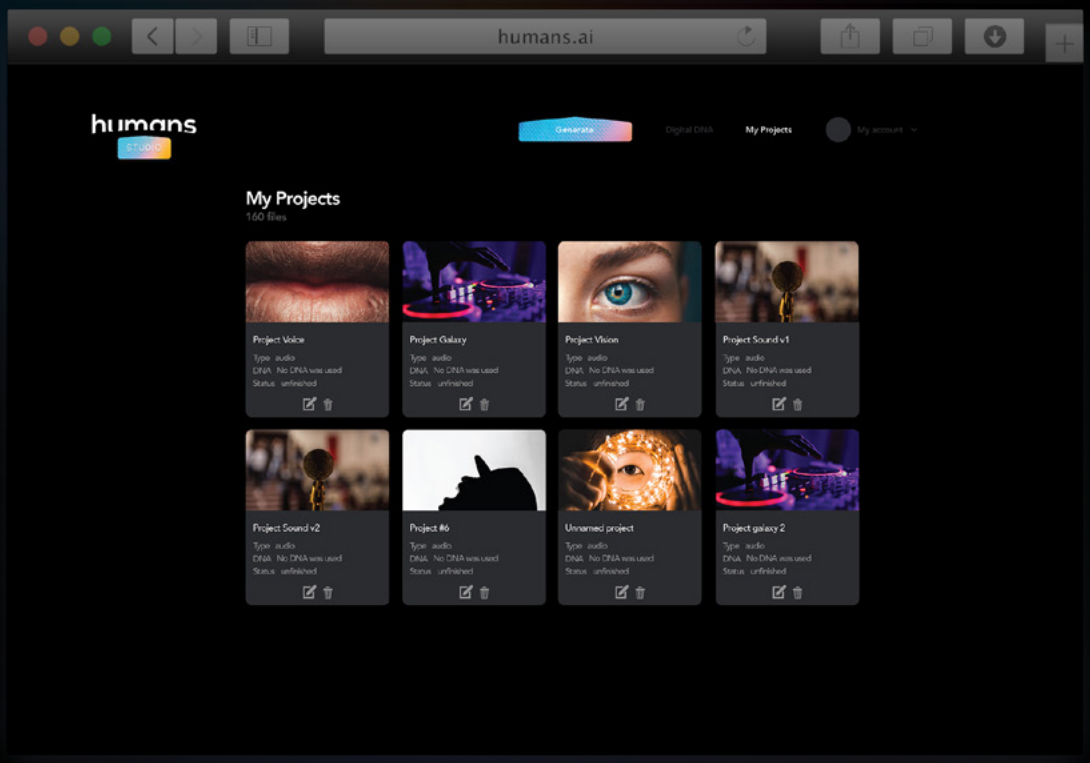

The Humans Studio is equipped with proprietary technology, allowing users with no prior knowledge of AI to play with recipes of combining neural networks to create ready-to-use results, such as AI-generated videos, audio recordings, etc. It will also allow software developers with no prior AI experience to combine recipes and add functionality on top to build apps (e.g., to create templates for AI-generated videos with preset after-effects). This way, the Humans platform will bridge the gap between AI researchers with no experience in building client-facing products and software developers with no AI knowledge and grow the AI community.

Currently, developers need deep AI knowledge to integrate multiple pieces of technology into the client-facing product. There is no solid framework for AI development yet, and building the product is almost like making it from scratch. Furthermore, getting access to the pieces that are not open-source might be cumbersome and take a lot of resources. Humans provides easy permissionless access to a variety of AI building blocks gathered in one place and the easy-to-use technology for seamless integration of these blocks. It will allow developers worldwide to use standard development frameworks and leverage Humans infrastructure to create AI-based products with no deep AI knowledge, thus jumpstarting AI adoption.

2.2.3. App Marketplace

Apps are combinations of AIs that generate valuable outputs on request. The request made to an App goes to each of its components that, in turn, share any proceeds respectively. So, for example, if the App consists of neural networks creating video and audio elements, each request will be sent to both neural networks and both will produce their part of the final result. Another type of combination is when the output of one neural network becomes an input of another one. So, for example, a neural network generating voice reading the text will produce an audio file which will be the input for a neural network producing video.

The Humans platform will collect the apps at its App Marketplace. Customers will be able to discover apps there to send their requests. For example, they can select the app generating audiobooks with the preset audio effects and background music and place the order to transform their text into the audiobook.

The marketplace will host an array of apps leveraging the underlying infrastructure of the Humans ecosystem. Now we are working with select partners on representative use cases to showcase what the Humans technology can do and provide examples for app builders. The first such project, Tovid.ai, is already built and live.

Independent tech developers can build more apps based on the framework available in Humans Studio:

From scratch, built to serve a specific or unique use case.

From recipes made available in the Humans Studio.

Users will have the ability to deploy apps at scale in the Humans decentralized ecosystem and access a variety of decentralized services provided by other community members. With time, some app builders may also choose to build their interfaces for their apps powered by Humans technology. So, generally, it will look like an independent service available via the third-party website (or offline) but using the Humans platform as their back-end infrastructure, in the same way that many online stores use Shopify technology. Alternatively, companies might build an app for their use only as a competitive advantage. In this case, they would also use the Humans platform as a back-end and would not publish their app in the App Marketplace.

2.2.3.1. Tovid.ai

While video content sells, it is costly to produce and impossible to be personalized. However, as the Humans' CEO, Sabin Dima, says, "whatever can be a video, will be a video". TikTok platform - hosting educational content on topics such as real estate or investments - is a vivid example. Moreover, advances in deep learning are enabling us to scale media production by generating synthetic media.

Tovid.ai allows transforming anything (from a newsletter to a website) into video through a SaaS platform that produces high-quality video content at scale using AI & Synthetic Media. The Humans' team developed the project to showcase the underlying technology. It is now live and onboarding B2B clients tapping into the market valued at $200B (video production, advertising and education).

2.2.4. PoH App Suite

The PoH (Proof-of-Human) App Suite represents the entry point of the Human Validator in the Ecosystem.

The human validator will share their biometric data, per industry and technology provisions on privacy requirements established by the European Union (currently the strictest worldwide).

Biometric data consists of personal data resulting from specific technical processing relating to a natural person's physical, physiological, or behavioural characteristics, confirming that this person is human, such as facial images or fingerprint data.

In the PoH App Suite, human validators will find the apps to perform their tasks:

(i) submit their biometric data to prove that they are human.

(ii) confirm that the AI settings were respected by providing a validation answer and attaching their PoH.

The human validator's action is twofold:

(i) It certifies that the AI has human oversight.

(ii) It certifies that the request processed by the AI complies with the rules of the relevant neural networks, data owners, and apps.

2.3. Token grid

The Humans ecosystem is based on the grid of the following types of tokens:

$HEART token

NFTs

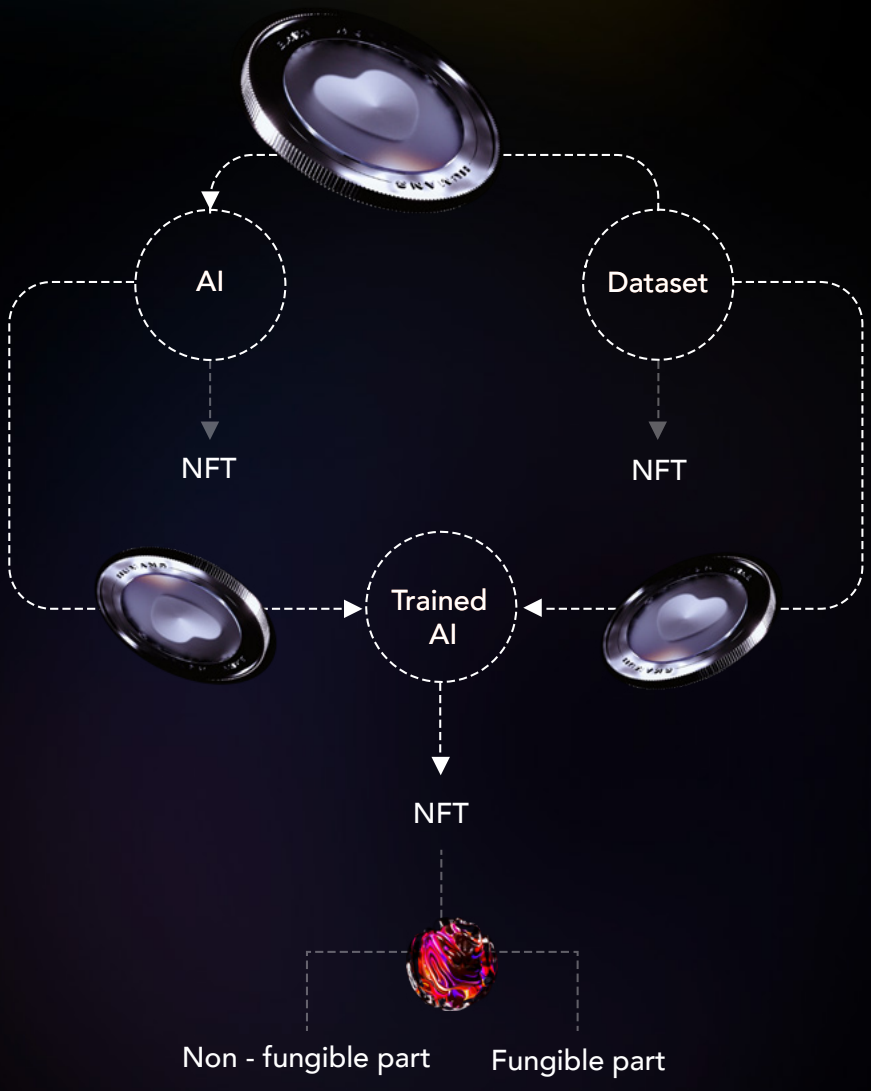

$HEART is the fundamental token of the Humans platform developed on the Humans native blockchain, powering the platform governance and all the payments. Each asset on the platform (AIs, Digital Genomes, trained AIs) has its own NFT. AI NFTs consist of a non-fungible part and a fungible governance token.

2.3.1. The $HEART Token

As the native token of the Humans ecosystem, the $HEART token empowers anybody to participate in the platform's governance and facilitates essential flows of value within it. All fees charged by the Humans platform will be paid in $HEART tokens, even when payment in the app originates in fiat or another accepted cryptocurrency. Generally, the app can choose what kinds of payments to accept from its clients, but it will still have to pay in $HEART tokens for the services provided at the Humans platform.

One of the primary functions of $HEART tokens is to facilitate staking. Staking $HEART tokens is required to perform key functions within the Humans ecosystem, including:

- Minting NFTs
- Account validation
- Becoming eligible as a Human Validator
- Creating apps
- Adding algorithms, AIs or Digital Genomes to the AI Library

Additionally, $HEART tokens will serve as the default incentive as various stakeholders contribute key value to the ecosystem. These incentives include rewards for:

- Creating a Tech provider account
- Uploading technology
- Creating the first AI NFT
- Participating in governance

2.3.2. Humans NFTs

Through NFTs, we can encapsulate digital information and maintain a persistent record of the value it produces over time. Humans utilizes NFTs to recognize assets within its ecosystem and ensure contributions to delivering valuable outcomes are adequately accounted for and recognized (so, every asset has its unique NFT). These assets include:

Algorithms: neural networks developed to perform certain AI functions with Digital Genomes and other data (before training).

Data: virtually any kind of dataset, in the case of synthetic media, it would be Genomes

(personal profile and biometric data, e.g. user's voice, likeness, etc.).

AIs: results of training Algorithms with Genomes or other Data.

Data providers combine AI NFTs into recipes to use for their purposes or to lease to customers for their projects. And Algorithm NFTs and Data NFTs are used to account for all the contributions to creating each AI NFT.

Whenever the asset is created or uploaded onto the platform, the author stakes $HEART tokens to add the asset to the AI Library, and the platform mints the new NFT. During the minting, the system will prompt the author to select settings regarding the NFT's governance and financial dynamics, including minting the fungible governance tokens for the NFT.

While the non-fungible part of the NFT represents the digital asset itself, the governance tokens represent the rights to govern the NFT. Governance tokens will be distributed across all the parties who (or whose data) contributed to creating an NFT. Only holders of the NFT's governance tokens can change the initial governance settings for this NFT.

Automated Liquidity Pools. The authors of NFTs (the ones who minted them) may act as Price Discovery Users. To do so, they may create a Uniswap-like liquidity pool to pair their governance tokens with $HEART tokens. It requires staking $HEART tokens against an equivalent amount of the NFT's governance tokens. So, generally, they can set any initial price for as long as they can lock the corresponding number of $HEART tokens.

Furthermore, if participants want to trade between two assets, a suite of smart contracts automatically facilitates purchases into and liquidations out of the NFTs' governance tokens and updates the exchange price in real-time, allowing for price discovery between the two assets.

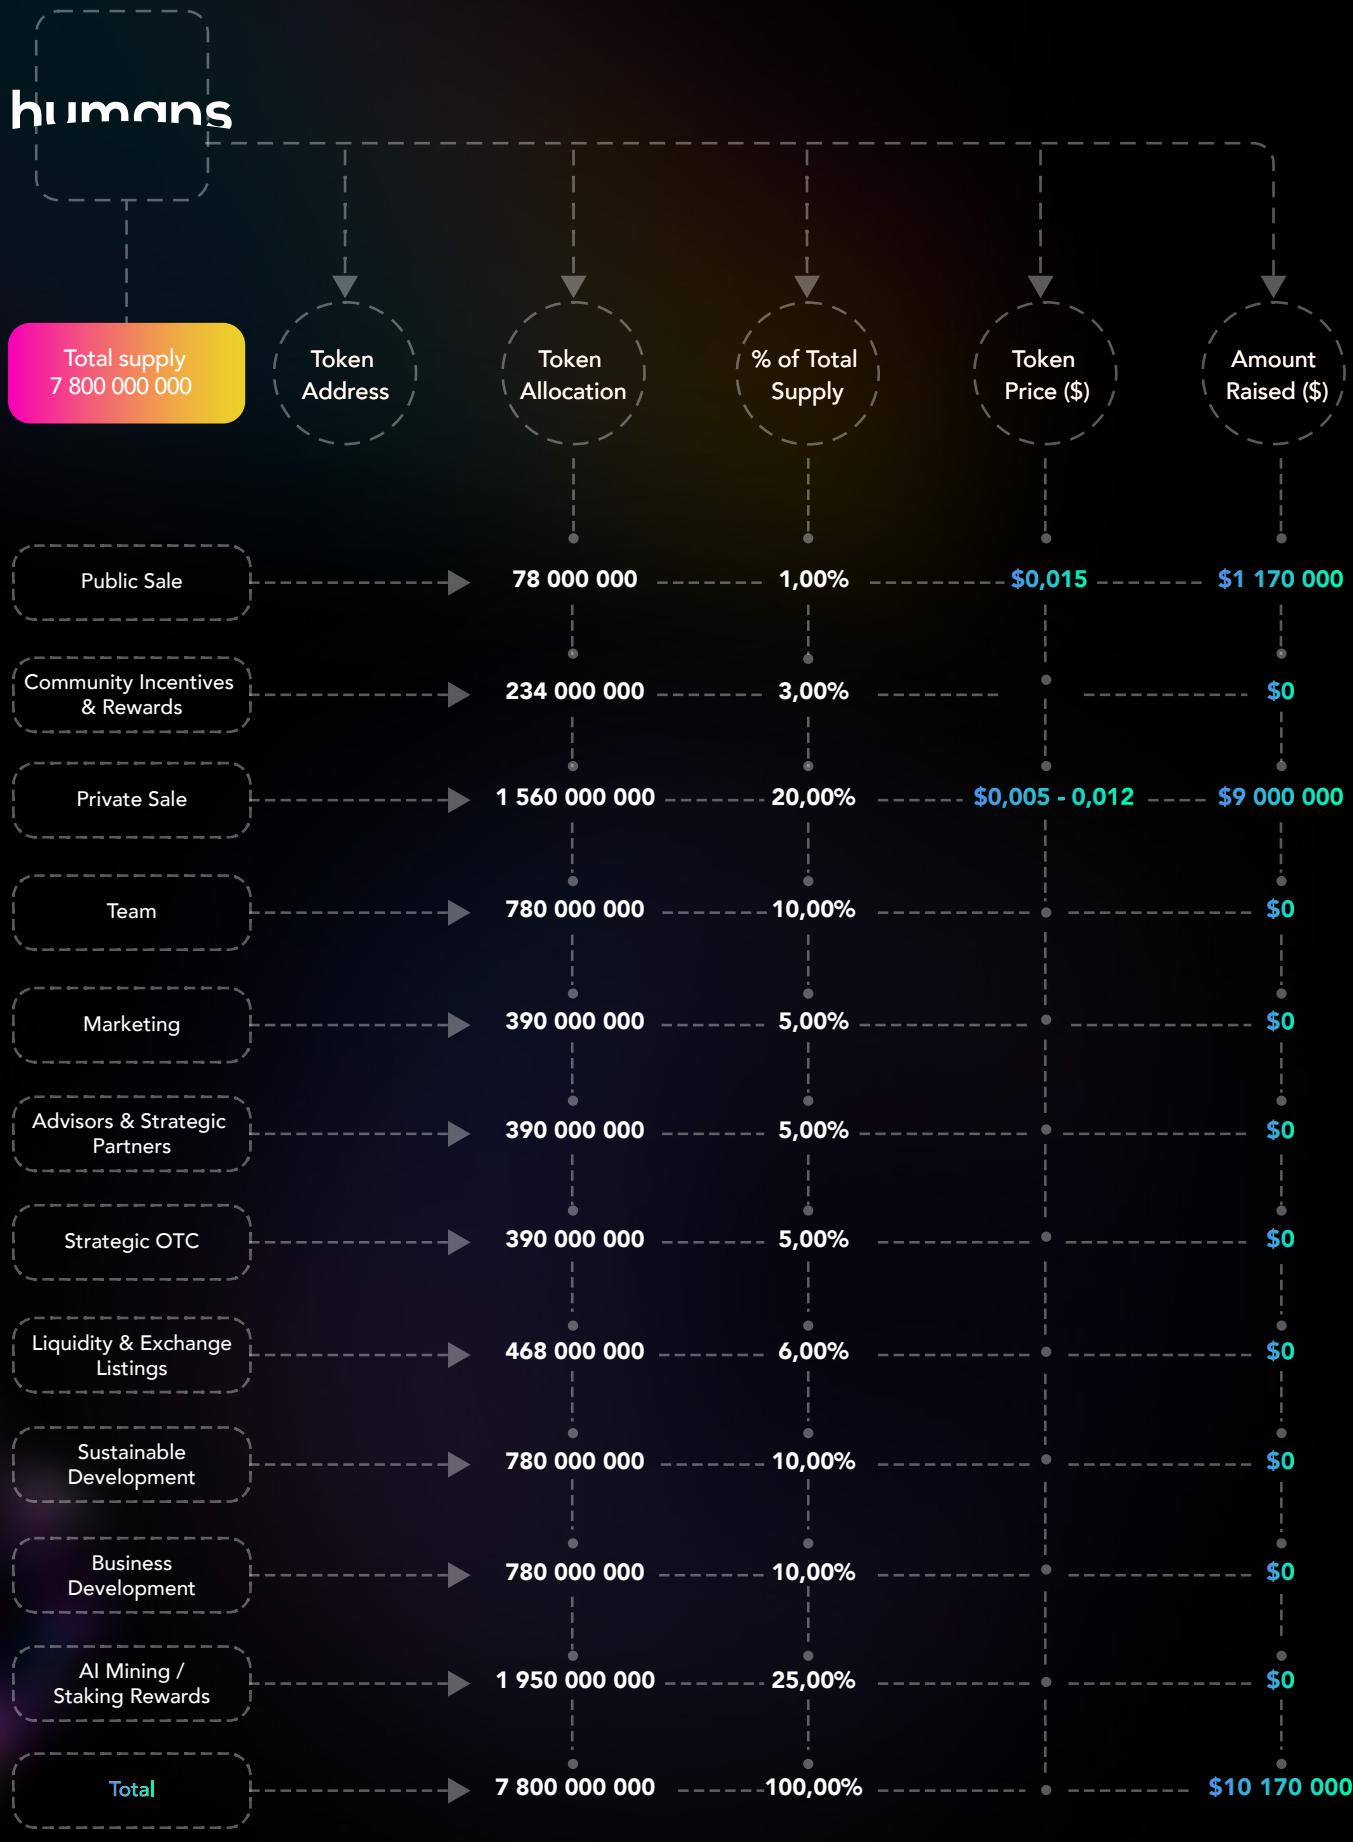

For a more detailed view, please refer to this spreadsheet on Humans Tokenomics.

2.3.3. Token minting and granting

The Humans blockchain relies on key stakeholders who contribute and ensure the good and fair use of the ecosystem. We can split the stakeholders into two main categories depending on their key function performed in the ecosystem:

Validators will mint new $HEART tokens as part of the consensus mechanisms:

Server validators

Human validators

Contributors will receive $HEART tokens as grants to fast-track ecosystem adoption in the beginning:

Tech providers

Data providers

App builders

2.3.3.1. Minting system

Validators are crucial entities in the Humans ecosystem and they perform two vital functions: security of the network and compliant use of the AI NFTs.

Server validators ensure the security of the platform by taking part in the PoS consensus mechanism and voting on minting new blocks to keep the system secure and working. They are awarded $HEART tokens for every new block signed on the blockchain.

Human validators ensure the good and fair use of the AI NFTs minted by the data providers and are further split into 2 main categories:

Agents

The requests are processed and split into subtasks which are sent for approval to the agents. They vote yes/no depending on the settings of the data product and mint tokens as part of the PoH consensus mechanism.

Block producers

They perform an essential role of actually validating the request after the agents validate the subtasks. They are called block producers because all the subtasks aregrouped into a block to be signed on-chain and essentially confirm the request. Upon realizing this, they are awarded $HEART tokens for mining the block and essentially validating the transaction part of the PoH consensus mechanism.

2.3.3.2. Grant system

To fast-track adoption, Humans will host growth campaigns to grant $HEART tokens to the community. The aim is for everyone to join and contribute to the ecosystem without having the token as a barrier or gatekeeper.

At the start, the tech providers will receive a grant in $HEART tokens when they make an account on the Humans platform, get it validated, and upload their tech. The grant is equal to the minimum stake needed to validate their algorithms.

The early adopters will be granted $HEART tokens for every step needed for the onboarding funnel. It is important to have as many users train their first neural network and expose their data product to the community to receive assistance with management and request validation. Not all early adopters will be crypto savvy, and having blockchain as the critical infrastructure and tokens as means of payments may create extra barriers for adoption. The grant system will help to remove these barriers. The grants will also be subject to certain conditions, which by design, will enable us to enhance each step to make the entire onboarding process as smooth as possible for the new entrants. Humans is in the unique position to showcase blockchain, crypto and AI in a straightforward solution.

All the granted $HEART tokens will have a locking period.

2.4 Economic Model

Tech providers, data providers and app builders maintain the co-production process resulting in new assets and corresponding AI NFTs. They, in turn, become the AI recipes that customers license. License payments made in any currency reach the Humans platform in $HEART tokens and are distributed between all the contributors.

Value Flow

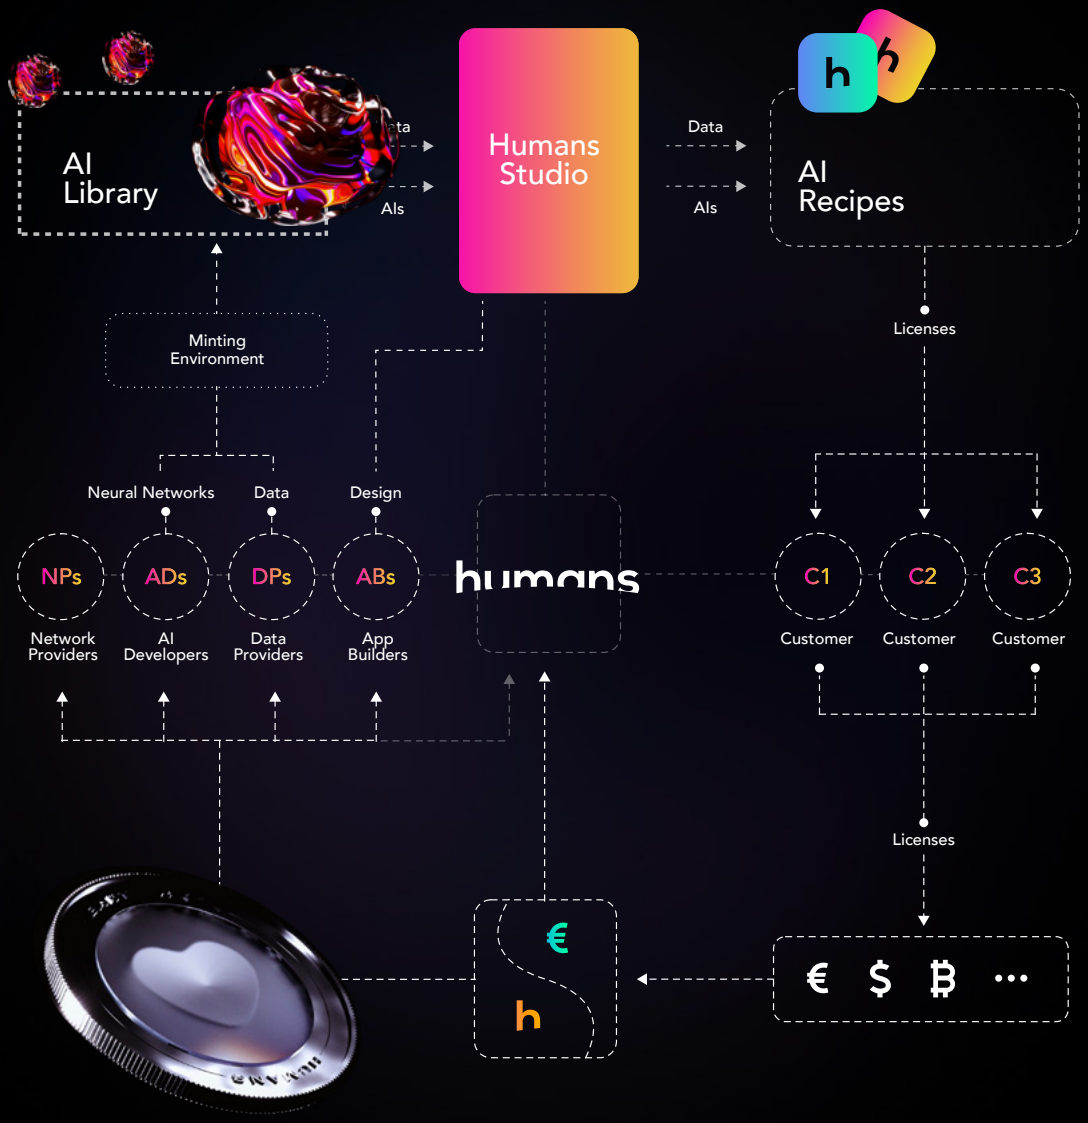

Part 1

CoProduction

NPs: Network Providers provide nodes to support the network
ADs: AI Developers provide AIs into the library
DPs: Data Providers provide Data into the library
ABs: App Builders design apps in the Humans Studio

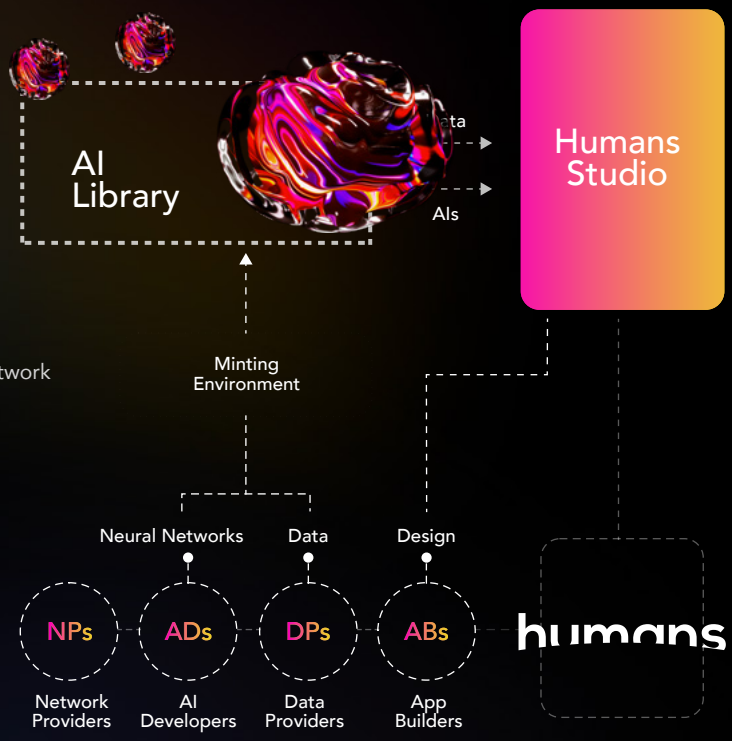

Part 2

AI as a Service

AIs and Apps are made available to customers via the Humans AI Marketplace; Pricing can be displayed in the Customer's currency of choice

Customers (C1-4) purchase use licenses (could also be labelled as Access); could also allow customers to offer bounties for specialty/bespoke AIs/Apps

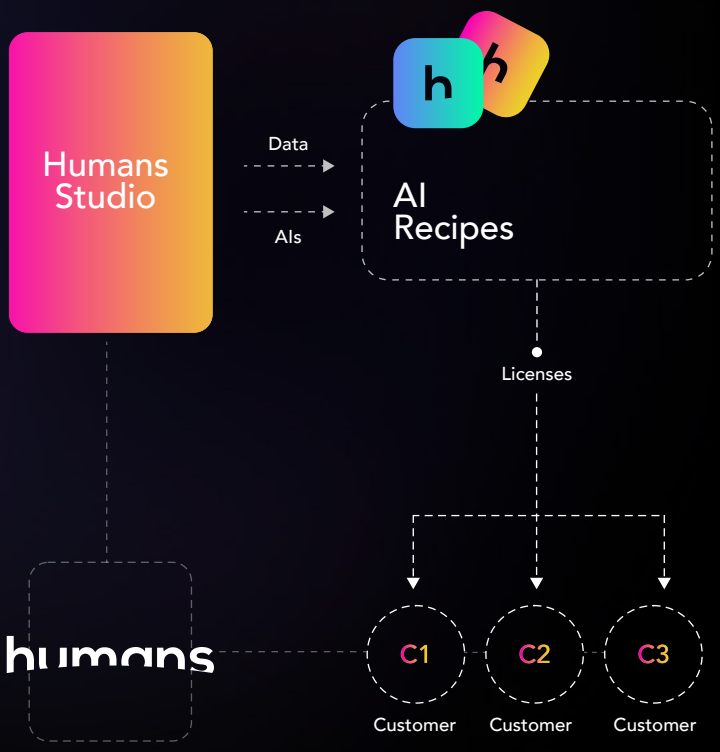

Part 3

Customer Payments

Customers can pay in any number of Currencies
Payments are automatically converted in $Human Tokens

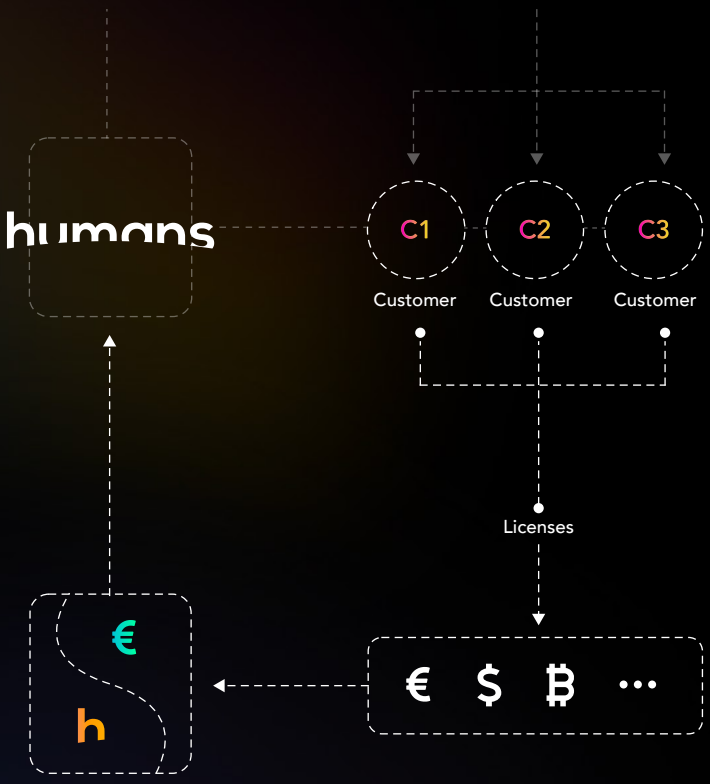

Part 4

Shared Earnings

All the parties that played a role in the creation of AI/App in question collects fees according to prior agreements determined based on the relative value of their contriabutions to the product.

Humans collects its own fees in the process.

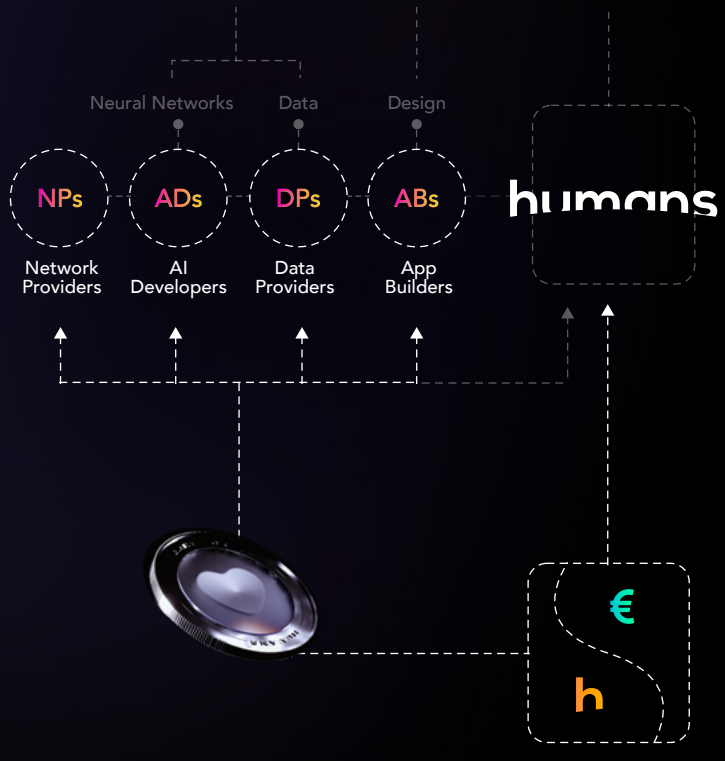

2.4.1 Fees

From the business perspective, we are building a marketplace, an AI of AIs, with the unique positioning that it is permissionless, and open to participation for the community to begin creating/contributing.

There are 2 types of fees to be paid in the Humans ecosystem:

(i) Fees paid in $HEART tokens to the AI NFTs:

The amount of $HEART tokens needed is determined by the AI NFTs' settings its author sets.

Owning the required amount of $HEART tokens is mandatory; otherwise, the transaction will not be processed at all.

(ii) Fees paid in the currency decided at the application level:

App builders will create apps that essentially are use cases/recipes of using different AI NFTs together. It will help to familiarize end users with both AI and blockchain tech in a way that is easy to understand.

App builders can choose both FIAT and cryptocurrencies as available payment types to make it easier for their customers to access their products:

- The apps solve the crypto barrier adoption by allowing customers to pay in FIAT as well;
- Important mention: the app will still pay all the fees for its request in $HEART tokens.

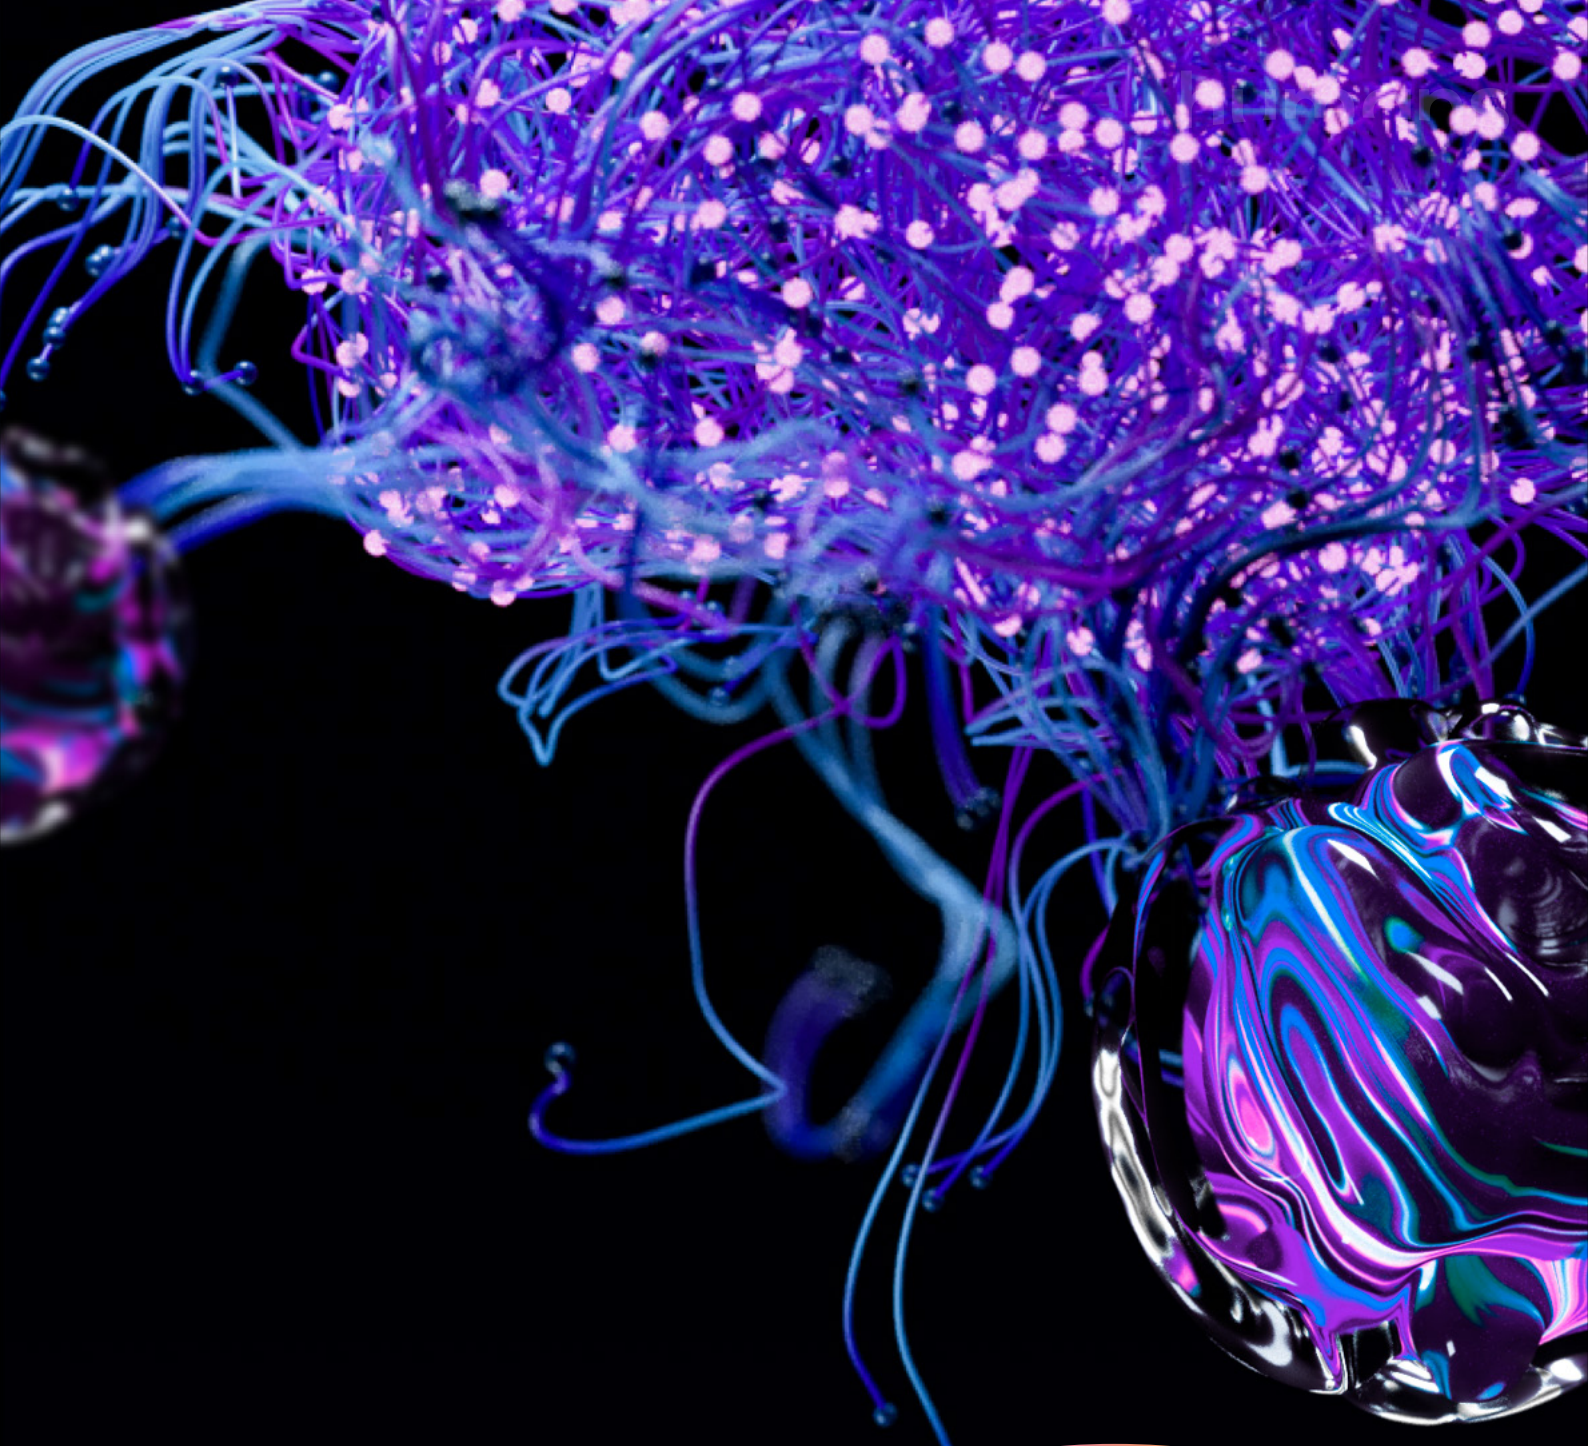

3

Humans Blockchain

3. Humans Blockchain

Humans Blockchain is envisioned as a Domain-Specific Chain built on top of current decentralized open-source technology, such as the Cosmos SDK. A decentralized open ecosystem powered by Human Validators, which stewards the good use of AI. Built with transparency at heart, a permissionless platform at scale in which anybody can create anything.

Humans technology implements modern infrastructure solutions with Cosmos powered by the Tendermint Byzantine Fault Tolerance (BFT) consensus algorithm as the base to optimize the build and address common blockchain problems. The array of ready-to-use instruments currently available on the market enables the straightforward development process.

For blockchain to enjoy widespread application, it is essential to interoperate and communicate with both external and internal blockchain protocols. It is a crucial aspect for our ecosystem, considering that the future apps built on top of Humans need to interconnect and interoperate with both decentralized and centralized systems.

Humans will build a domain-specific chain with multiple layers of consensus, supporting both PoS- and PoA-based consensus mechanisms. Thus, it will allow us to experiment across different infrastructure layers regarding the best consensus and interactions depending on the job for the specific task.

Blockchain development will also need to take into consideration composability at the consensus level to ensure each application has its sovereign execution space.

In practice, it means that the AI NFTs in Humans can have different consensus rules depending on the settings of the specific AI NFT, app type, reputation, etc.

We can leverage different mechanics around when the computation can be performed in the life cycle of a request and when the consensus comes into play. Thanks to it, we can fast track adoption for centralized applications and systems, which is crucial because a substantial part of the value will be derived from regular businesses.

To encapsulate all this into a compelling tech solution to be adopted at scale, everything needs to come under one big umbrella, as in a domain-specific chain.

3.1. Domain–Specific Chain

The Humans Ecosystem is designed to contain and manage AI at scale. We are building a Domain-Specific Chain to accommodate and create synergy between all the main stakeholders:

- AI researchers
- App builders
- Data providers
- Human validators
- Server validators

Through Humans, anybody can be a creator, participate in governance and validate the AI decisions (part of Proof-of-Human consensus). The Humans Blockchain empowers them to join and build on top of Humans:

Create – empower the user through AI and blockchain to be in control of their data and creations

---

Decide – the users are incentivized to participate in governance to decide on essential aspects regarding the future of the respective AI

Earn – the user gets rewarded proportionally with the usage generated by their AIs

We use the Cosmos SDK framework to create an environment tailored to raise and steward AI.

Cosmos is the first framework providing an interchain operability solution that is viable and currently functional. Cosmos is a very robust open-source framework, which allows for the deployment of sovereign blockchains. Cosmos secures over $170 bln in digital assets across more than 40 apps and services. It has a proven track record, having been chosen by large projects like Terra, Osmosis and Thorchain.

The Regen Network is the most representative project so far to build on top of the Cosmos SDK and launch a Domain-Specific Chain. They have shown that this framework can be used to create a new decentralized domain-based economy with intrinsic rules and values tailored explicitly for their community.

Main advantages of building a domain-specific chain on top of the Cosmos SDK:

The state machine is connected to the underlying consensus engine through the Cosmos' protocol ABCI (Application Blockchain Interface)

With the newly in the works ABCI ++, it is possible to separate the consensus engine from the compute engine

Programmable consensus, starting with Tendermint

Flexibility on the programming languages supported and developers' support

---

The ability to tweak the framework and consensus

The ability to implement automatic execution of code

The beauty of building on Cosmos is that we inherit several pieces of tech already validated by the use in production by several other notable projects like Rune, Terra, and others:

Tendermint BFT

Our tech architecture will have several layers of consensus depending on the task at hand to be resolved.

We will inherit the Tendermint BFT for our layer 1 DPoS consensus layer, which ensures the platform's security.

For the layer-2 solution, requests get validated by HV block producers through PoH, which runs a modified version of Tendermint BFT to reach consensus around validating a request.

Interchain

We are building a decentralized network consisting of AI DAO NFTs, powered by the Proof-of-Human consensus algorithm. This is envisioned to tailor to the needs of HV block producers to accommodate the signing (validating) of blocks (requests).

For a widespread application, the inter-blockchain communication is essential to empower the apps to be built on top of Humans to interconnect and interoperate with both decentralized and centralized systems.

PoA

It will allow us to experiment across different infrastructure layers to accommodate big enterprise companies that will run the execution on PoA sidechains to ensure that the state on-chain reflects the reality.

Composability

Composability at the consensus level is necessary to separate the decision on accepting a request by an AI NFT from actually performing the computation needed to generate the output solicited.

3.2. Governance

Governance is paramount in our ecosystem as we are putting humans at the forefront of the AI revolution with the sole objective to align the AI's objective with ours. Governance stems from everything we do and it will play a central role in our ecosystem.

We have governance at two different levels: protocol level and AI NFT level. The protocol level governance is well documented and will respect the current governance standards and workings with communities like Cosmos, Osmosis, etc.

Protocol level

Concerns with the control of a number of parameters by the protocol and these parameters apply to all AI NFTs.

NFT level

This concerns the control of parameters that apply to a specific AI NFT.

In the case that the owner of the NFT opened the data product to the community, then the control is in the hands of people that own governance tokens for the respective AI NFT (the fungible part of the AI NFT).

The protocol level parameters will serve as defaults for any newly minted NFT and can be changed by governance at the NFT level.

Examples of such parameters that can be overwritten at the NFT level:

- Modify standard governance proposal time set at the protocol level;
- Modify Approval/Rejection timeframes for requests sent to the AI NFT;
- General types of requests accepted.

Role of human validators as DAO participants:

Manage

Approve/reject request sent to AI NFT via Proof-Of-Human

Govern

Override certain Global Governance via the proposal submission

We are proposing a new way to govern AI NFTs like DAOs in which the owner can choose to share a part or all of the control with relevant community members who may want a stake in the respective data product. Considering the AI NFT will generate economic value in direct proportion with the usefulness of the NFT, the owners of the AI NFT DAO tokens will have a say over the pool of $HEART tokens collected by the AI NFT from accepted requests. For example, the data provider of voice can set restrictions on what kind of requests for voice minting can be accepted (e.g., only business videos), and human validators will check every request to comply with these restrictions. However, if the data provider chooses a DAO style for the product governance, the community members will have a say on what restrictions to impose.

Governance Flow

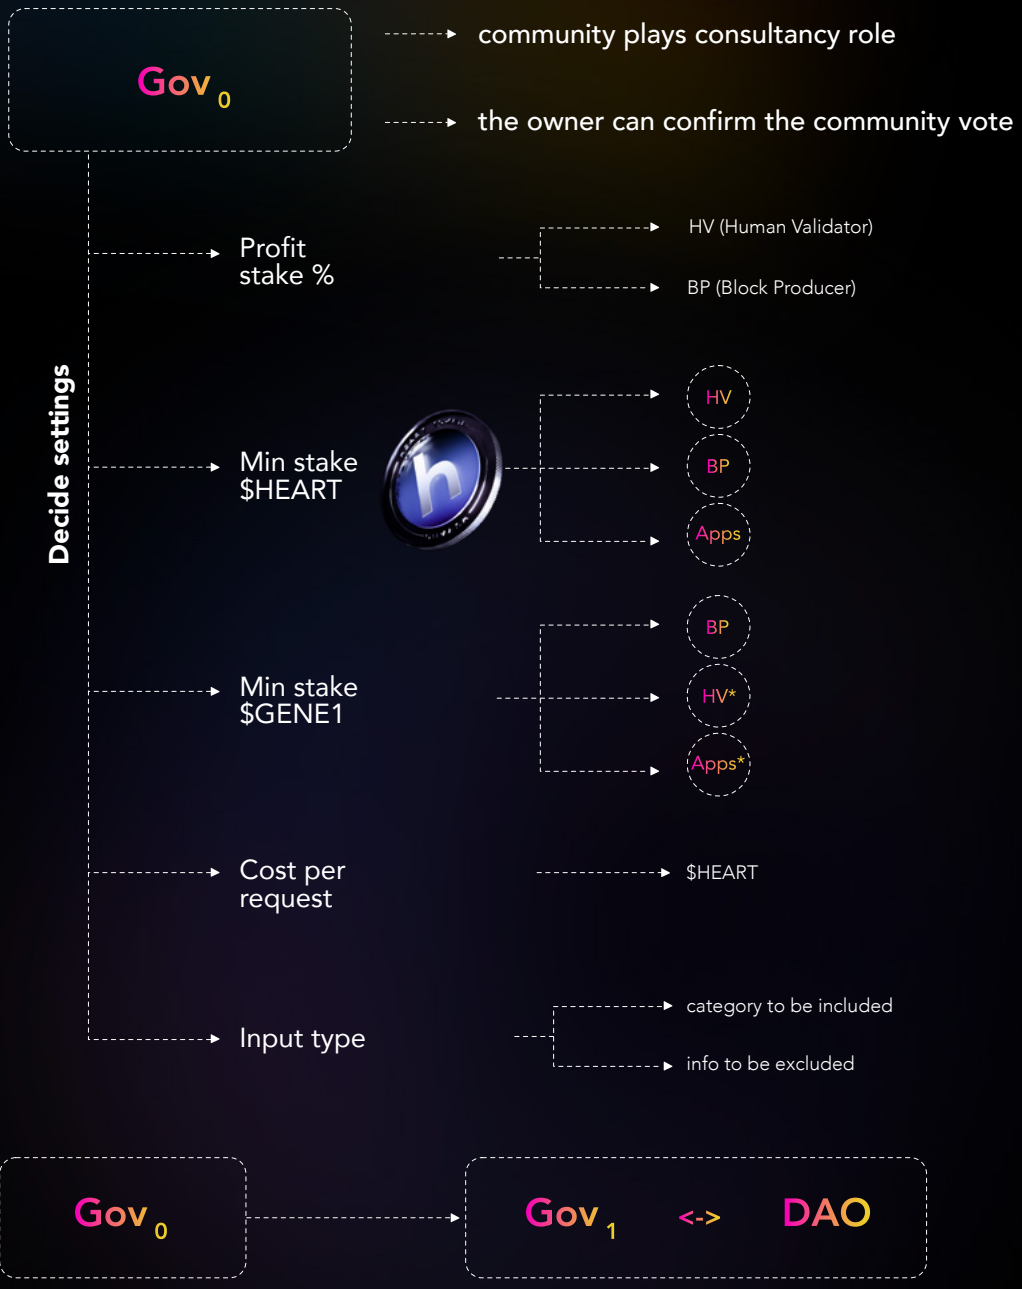

Gov 0 ⟶ Gov 1 <-> DAO

Owner switches to DAO ⟶ the community's decision is final

3.3. Validators

Validators play a crucial role in the Humans ecosystem. We can split them into two main categories depending on the specific tasks they perform to ensure that the blockchain is up and running and the data products are used according to the rules.

3.3.1. Server Validators (SV)

The Humans platform prioritizes the accessibility of AI. To fulfil this goal, we needed to create a solution for permissionless access to the technology. Thus, the idea of decentralization via server validators was born. Complementing this with a staking mechanism aligns the interests of validators with those of data providers and tech developers. It ensures that AI works within the prescribed parameters. Generally, blockchain appears to be the simplest way to crowdsource AI governance.

The role of server validators is to run a full node, produce blocks, validate and participate in the Delegated-Proof-of-Stake (DPoS) consensus. They ensure the network's security and provide decentralization, participating in the DPoS consensus mechanism by transmitting votes (cryptographic signatures) to agree upon the next block. Server validators are responsible for validating transactions, and ensuring all parties contributing to a valuable outcome receive appropriate compensation. Their voting rights are determined by the amount of $HEART tokens pledged as security.

3.3.2. Human Validators (HV)

The Human Validators are the primary agents who ensure that only the valid requests received by a data product are accepted and processed. They participate in the Proof-of-Human consensus mechanism by sharing their biometrics. Their role in the network depends on the amount of the staked $HEART tokens and data product fungible tokens. Depending on these conditions, we can split the human validators into two categories depending on the function they serve in the system:

Agents

The requests are processed and split into subtasks which are sent for approval to the agents. They vote yes/no depending on the settings of the data product. For example, a text from a request to generate a minted voice can be split into sentences and sent to different agents to validate compliance with rules. Their aggregated votes from all the subtasks are sent for final approval to the block producers.

Requirements:

Satisfy the minimum staking requirement imposed by the network for $HEART

Additionally, each AI NFT (trained AI algorithms) can set up different minimum staking requirements and additional rules

Block producers

They perform a vital role of actually validating the request after the agents validate the subtasks. They are called block producers because all the subtasks are grouped in a block to be signed on-chain and essentially confirm the request.

Requirements:

Satisfy the minimum staking requirement imposed by the network for $HEART

Satisfy the minimum staking requirement imposed by the AI NFT for the $GENE1 (fungible part of the AI NFT)

3.3.2.1. The onboarding process for HV

1. Download PoH App and create an account.

2. Set up an agent profile and select how the PoH will be performed.

3. In-App onboarding simulating subtasks:

Agents will go through gamified training (HV training) to perform a few tasks and earn $HEART tokens as prizes. This training is meant to onboard agents and inform them about performance requirements.

4. Stake $HEART tokens:

Agents will receive the minimum stake required by the system by completing the HV training. Considering that each data product will have a required minimum stake, agents may need additional stakes to validate for as many data products as possible.

5. Become a block producer (optional):

To upgrade their status and tasks performed, agents can buy the additional data products' fungible tokens and stake them to claim higher rewards from the data products they validate the most.

3.4 Consensus

3.4.1. Proof-of-Human (PoH)

3.4.1.1. Overview

Proof-of-Human is the cornerstone of our decentralized ecosystem. Everything we do stems from the fact that we put humans at the forefront of this AI revolution. Humans become validators in an attempt to build the framework to contain, manage and govern AI.

Proof-of-Human is a new blockchain consensus that relies on humans to leverage their biometric data to ensure close biological supervision of every AI. Humans empowers people to participate in governance, management and deciding the fate of any AI, essentially making sure that the objective of the AI is aligned with the human's objective.

3.4.1.2. Request Flow

Apps on top of the Humans ecosystem are a collection of multiple AI recipes and represent the main point of contact.

1. The app generates a request containing the input to be processed and the number of $HEART tokens required by the AI NFT.

2. The request is inserted in the mempool, and before the processing starts, the system checks if the amount of $HEART tokens sent meets the requirement of the AI NFT.

3. The input to be approved is processed into representative sub-tasks to be sent to the HV Agents who respect the minimum staking requirements:

STAGE 1
The input is chunked into subtasks;

STAGE 2
Assert which subtasks need validation and create the validation mesh;

STAGE 3
Optimise and batch sub-tasks to be sent to the HV agents for approval.

---

4. The sub-tasks are sent to HV agents for approval:

Each sub-task is sent to N agents and the aggregated response is collected and sent to block producers for final validation.

---

5. Block producers receive the request input in the form of a block with multiple transactions (each sub-task is a transaction), and they need to reach a consensus if this block should be signed on the blockchain or not.

The block producers are a known group of human validators, each identified by their biological control key. Every request (block) they attempt to reach a consensus on one request at a time, where the block is the list of subtasks sent to HV agents.

During a block validation interval, a block producers' set is defined as the set of validators signing transactions that agree to confirm the next request (sign the block). This validators' set is dynamic and changes as validators join or leave the consensus process (they stake or unstake the required $HEART tokens and $GENE1 tokens).

Furthermore, voting for consensus on every request proceeds in rounds, and every round has a leader who proposes a block. Block producers then vote on whether to accept the proposed resolution for the request or move on to the next round. The proposing block producer for the round is chosen deterministically from the ordered list of block producers, in proportion to their voting rights, determined by their stake of $HEART tokens, $GENE1 tokens, and other factors, like reputation (depending on the AI NFT settings).

Enforcement of the AI NFT settings will be assured by using optimal Byzantine fault-tolerance via supermajority voting and the staking mechanism mentioned above. For example, suppose any set of HV agents or block producers ever attempt or even succeed in disrupting the request approval flow. In that case, the protocol will identify and slash them (they will lose all their tokens or a part of them depending on the violation).

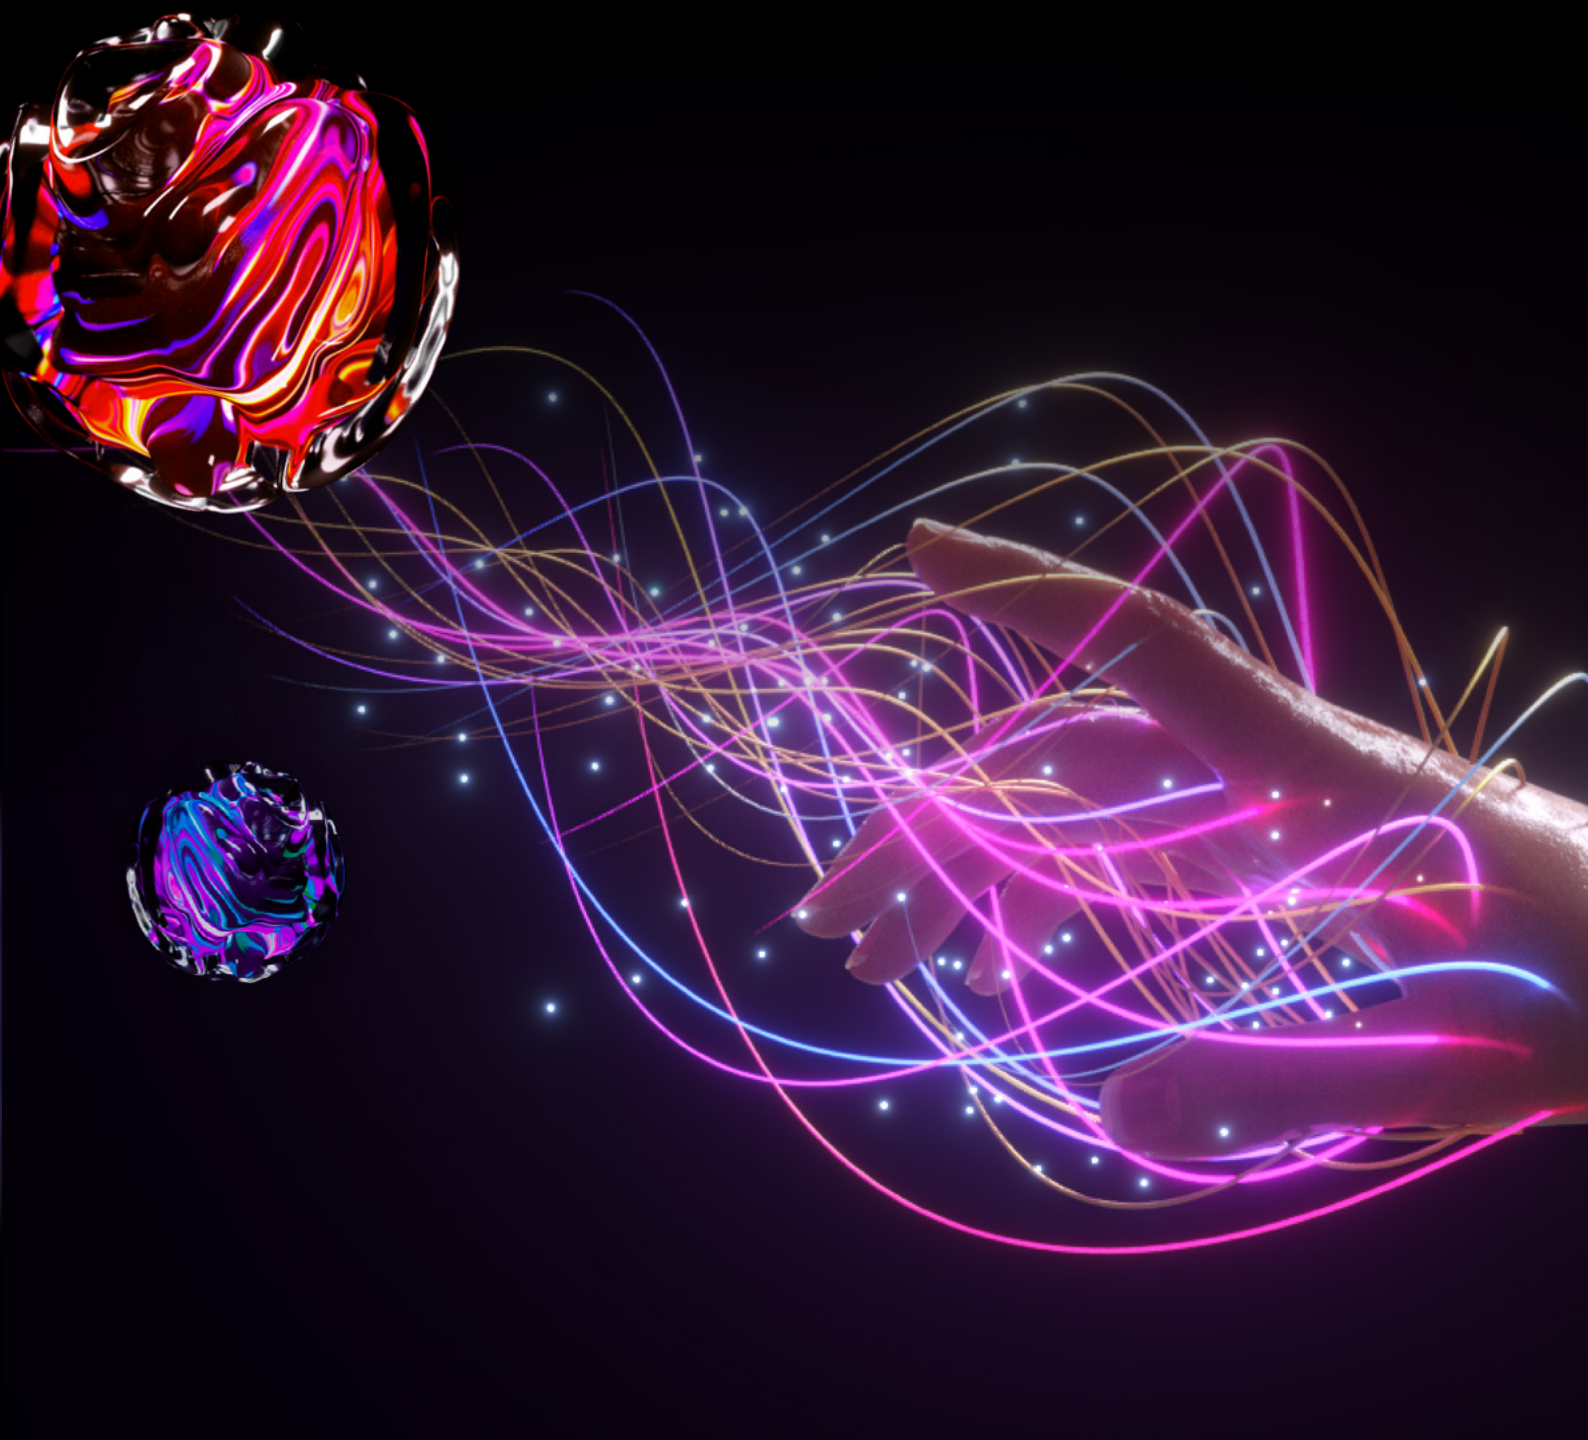

POH Flow

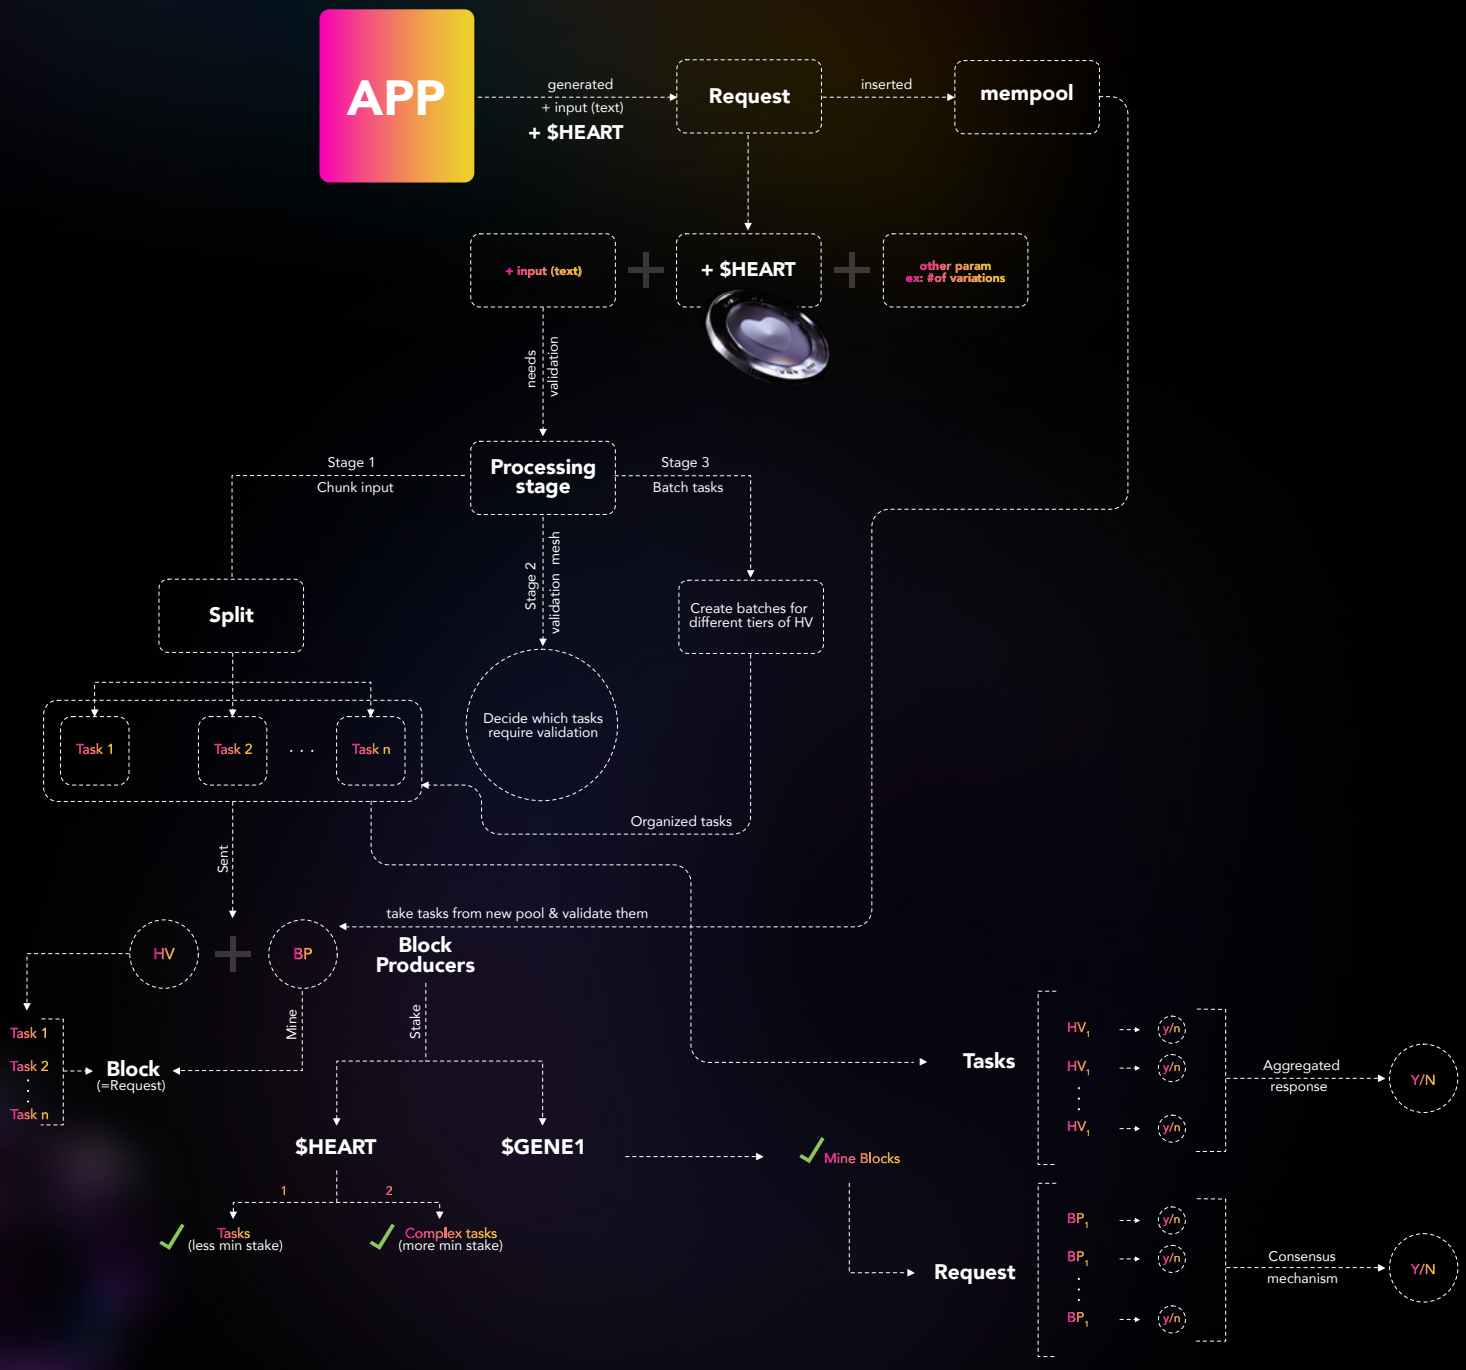

For a more detailed view, please refer to (image).

3.4.2. Delegated-Proof-of-Stake (DPoS)

To ensure the security of the network, we will use a BFT-based consensus mechanism. Server validators validate transactions and add new blocks to the blockchain. They participate in the DPoS consensus protocol by broadcasting cryptographic signatures that act as votes to extend the blockchain.

The SV must lock up a predetermined number of $HEART tokens to become a server validator and have positive voting power. This can either be self-funded or acquire the voting power from other staking token holders by having them delegate stake to server validators.

The delegators commit $HEART tokens to contribute voting power to a validator of their choice to earn a portion of the block reward. Rewards are distributed to all stakers proportionally to their stake at every checkpoint, except the proposer getting an additional bonus.

On the other hand, validating does have some risks. Stakes are at risk of getting slashed if the validator node commits a malicious act like double signing or validator downtime affecting the linked delegators at that checkpoint. Delegators are putting their $HEART tokens at stake and may lose these tokens depending on whether or not the SV behaves as intended by the ecosystem.

As long as 2/3 of the weighted stake of the validators is honest, the chain will progress accurately. Validators stake their $HEART tokens as collateral to work for the network's security and earn more $HEART tokens as a reward in exchange for their service.

3.4.3. Proof-of-Authority (PoA)

We are looking to use PoA for the AI NFTs which remain under the full management of the data provider (we will call this "centralized") and for specialized AI NFTs down the road (e.g., health/medical related AI NFTs).

The data provider can choose the specific HV block producers to confirm the transactions. It excludes most participants and creates a situation where only a selected few are responsible for ensuring the requests received by the AI NFT are validated correctly. For example, for health/medical related AI NFTs, data providers would choose HV block producers with a medical degree or background.

The select few, in most cases, will represent the data provider itself or an appointed entity that become/s the designated server validator. This is highly necessary in cases where trust is implicit for high-performance apps which work with specific AI NFTs (which they either own or have a business relation with).

To ensure that the underlying neural networks are used in a compliant way, we will log periodically on the chain the usage of a specific AI NFT. We will create a service responsible for checkpointing to aggregate blocks into a single merkle root and periodically publishing it to the Humans blockchain (further tech considerations need to be addressed in the publishing process).

3.5. Ledger Functionality

3.5.1. Human POS Protocol

The Human Proof-of-Stake (PoS) protocol is a form of consensus mechanism that works by selecting validators in proportion to their quantity of token holdings. This system was chosen to avoid exorbitant energy consumption which the Proof-of-Work (PoW) protocol would involve.

For a transaction to be approved, it must be appended to the blockchain. Validators accomplish this appending and earn a reward for doing so. For the blockchain to remain secure, it has a mechanism to prevent a malicious user from taking over a majority of validation. We accomplish this by requiring validators to have a specific quantity of tokens, requiring potential attackers to obtain a significant fraction of the tokens on the blockchain to carry out an attack.

The protocol can suffer from the nothing-at-stake issue, where validator nodes confirm contradictory copies of the blockchain because there is an insignificant cost of doing so and a smaller risk of losing out on compensations by validating a block on the opposite chain, allowing for double-spending. This is mitigated by penalising validators who confirm conflicting chains and by structuring the rewards so that there is no financial incentive to generate disputes.

3.5.2. Identity, key management and arbitration

For the parts of the system to function, the Humans blockchain compute engine will have to resolve matters regarding identity, key management, and arbitration.

The identity is continuously under the administration of its owner. Humans will implement the design of an existing portable digital identity and key management solution, such as the Sovrin Network or similar.

As smart contracts cannot yet encode certain rights and privileges that require subjective constraints, Humans will use a decentralized oracle protocol, such as Aragon or similar, to mediate subjective disputes.

3.6. Nodes

3.6.1. Compute nodes

Compute nodes are virtual machines running distributively across many computers that handle all types of compute requests, from oracles to training AI Models and running AI networks.

It is necessary to understand containerization as a type of virtualization; we will look into how we can distinguish a virtual machine from a container by focusing on characteristics in contrast.

Virtualization is a technique that provides fully functional computing power without a demand for resources, physical location, proximity or organization. This technology allows us to share and allocate resources within multiple executing environments. Without context, this might sound like a vague term, but in our ecosystem, it encapsulates the abstraction of network servers environments and different parties from their concrete hardware requirements, thus allowing us more flexibility, more functionality and a significant impact on the resources available.

Portability and encapsulation has been added as a general benefit: the unique ability to migrate resources live from one point to another without suffering from an outage. This has been made possible by leveraging two free and open-source technologies, KVM and QUEMU.

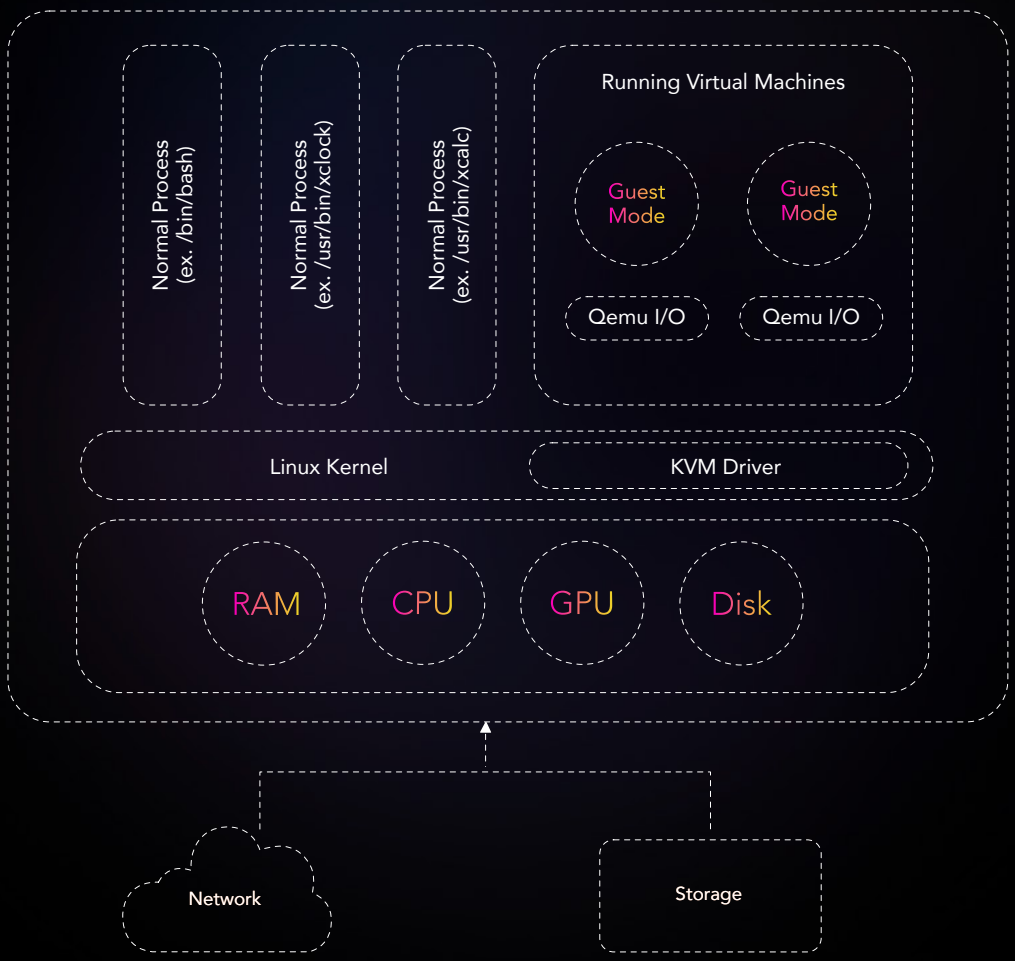

3.6.2. Validation nodes

The Humans ecosystem is based on the Proof-of-Stake (PoS) mechanism, which relies on validators to secure the network.

The role of the validators is to run a full validation node and participate in the consensus by broadcasting votes which contain cryptographic marks signed by their respective owners. Validation nodes commit new blocks in the blockchain and receive rewards in exchange for this to happen. They also have a role in the governance by expressing votes on new proposals, and these votes are weighted according to their total stake.

Staking is the process where the validator's weight is determined by the amount of staking tokens ($HEART) locked as collateral.

These $HEART tokens can be staked directly by the validator. After which they become validators, and the system will calculate how often this validator will have to propose a block and the reward for the proposal.

Validators are required to run a full node; a full node is a program that validates the transaction and the blocks of a blockchain. It is distinct from the other nodes in that the other nodes only process block headers and only a subset of the transactions, thus requiring the full node to consume more resources to be a validator.

In practice, running a full node means getting the software to run on a computer with low network latency and without downtime.

In the following diagram, we explain the architectural differences between a traditional blockchain using proof of work and our blockchain using proof of stake.

4

System Architecture

4. System architecture

4.1. High-level architecture

This high-level architecture provides the architectural view of the Humans ecosystem and its components. It has been developed to support the full range of the Humans ecosystem functionality.

The visualized overview can be found here.

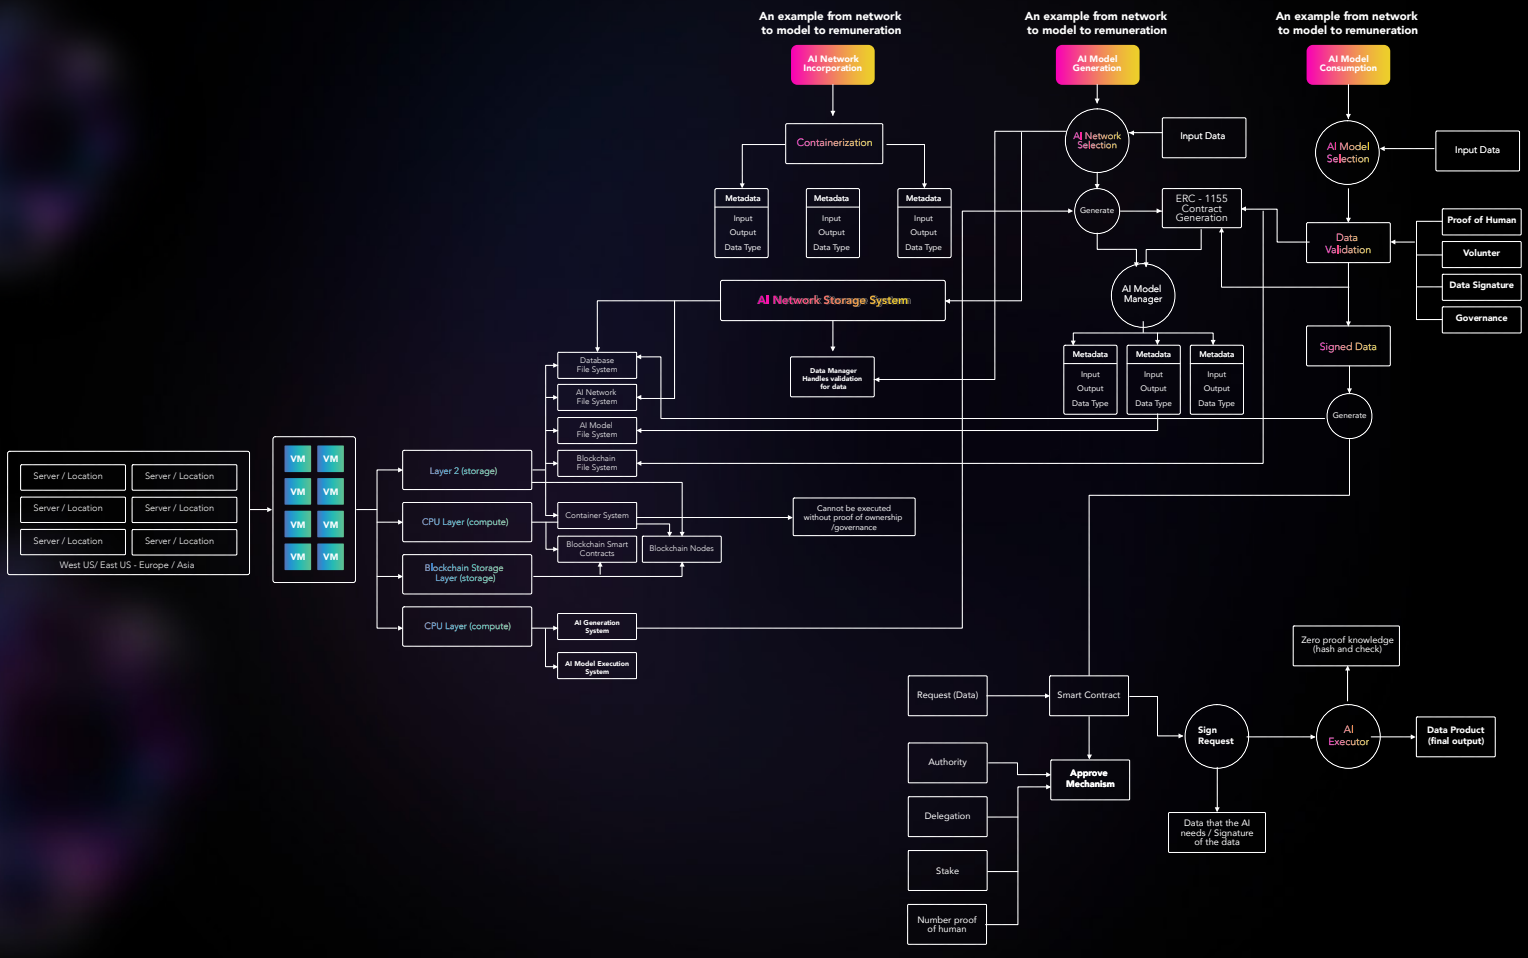

4.2. Humans Network Components

Humans network components are a specification of libraries and plugins to configure network interfaces inside the ecosystem for compute engines. Network components are responsible only for the network connectivity of compute engines and removing allocated resources when the compute engine is deleted.

The chosen architecture achieves private networking (communicating not over the internet, but rather through a secure private network), snooping tampering (traffic is encrypted point-to-point), no third-party reliance to secure traffic, and compression (faster data transfer).

4.2.1. Ingress Network Engine

Humans ingress engine is an API that provides routing rules to manage external 3rd party access to the compute engines inside the ecosystem.

We can easily set up rules for routing traffic without creating a bunch of Load Balancers or exposing each service on the compute engine node. It can also be configured to give compute engines externally-reachable URLs, load balance traffic, terminate SSL / TLS, and offer name-based virtual hosting. Finally, this ingress does not expose arbitrary ports or protocols.

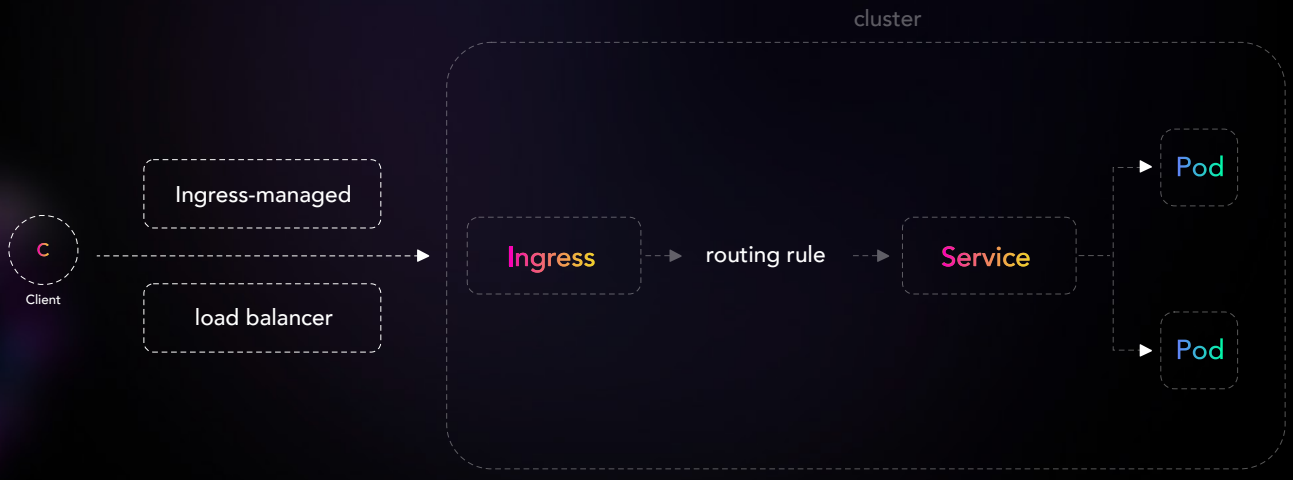

4.2.2. Egress Network Engine

Egress network engine in our ecosystem implies traffic existing outside an entity or a network boundary, while the ingress network engine enters the network boundary.

Since traffic is often translated using NAT (Network Address Translation) in and out of our private networks, a response from a public endpoint to a request initiated inside that network is not considered Ingress.

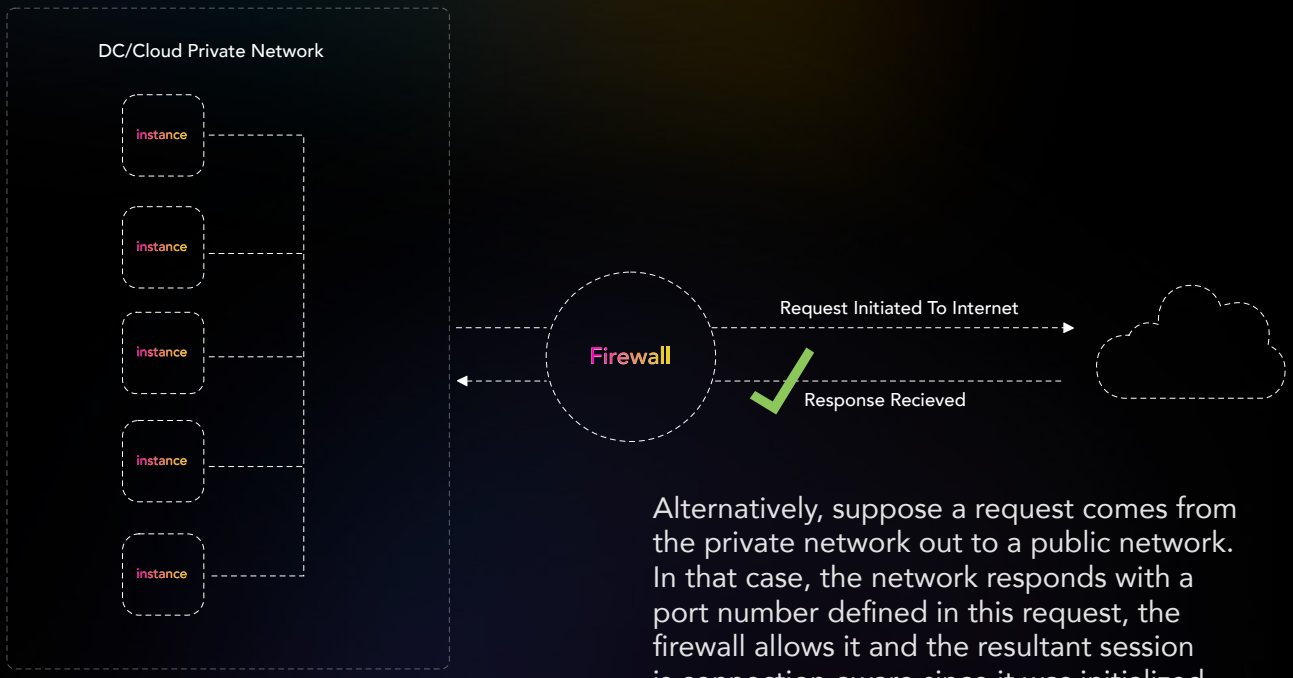

Alternatively, suppose a request comes from the private network out to a public network. In that case, the network responds with a port number defined in this request, the firewall allows it and the resultant session is connection-aware since it was initialized based on the supplied port number. Following this, the traffic is considered egress.

4.2.3. Service Discovery Engine

The service discovery engine is responsible for automatically detecting devices and services inside the ecosystem. Service discovery protocols are a networking standard detecting networks by identifying resourcewvs. In the ecosystem, the service discovery engine helps reduce configuration efforts of compute engines presented with compatible resources such as a specific AI network or a specific AI model. We have extended this concept ecosystem-wide at the compute engine level to discover and access compute engine components.

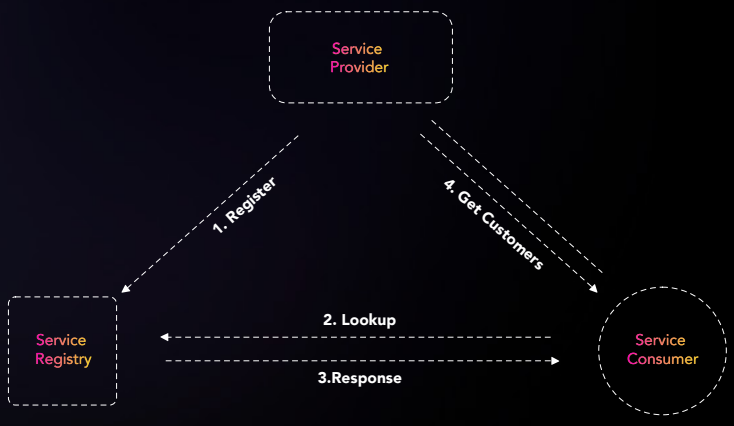

The service discovery engine can locate a network automatically, removing the need for a long configuration process. It works by compute nodes and components connecting through a common language on the network, allowing devices or compute components to connect without manual intervention.

The ecosystem uses two types of service discovery: server-side and client-side.

Server-side service discovery allows clients applications to find services through a router or a load balancer. Client-side service discovery allows clients' applications to find services by looking through or querying a compute engine service in which compute engine services are all within.

4.2.4. Load Balancing Engine

Load balancing plays an essential role in the Humans ecosystem. This engine is responsible for distributing network traffic across multiple compute engines. It ensures that no compute engine or compute node single-handedly bears too much demand. By spreading the requests evenly, we improve the application responsiveness. It also increases the availability of compute engines and apps for 3rd parties.

Besides the traditional algorithms, we load-balance requests towards the compute nodes based on different circumstances and metrics, such as uptime, technology, performance, load, security and location. It's a stepping stone in the full decentralization of our compute engine nodes.

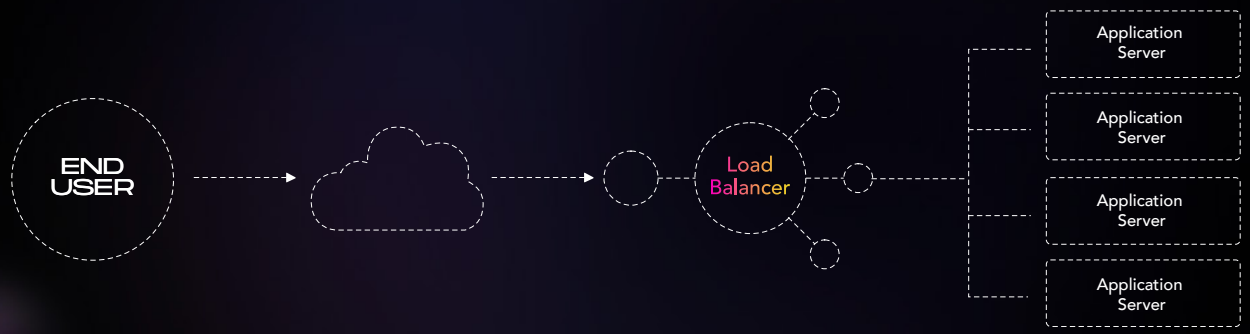

The load balancing engine also plays a vital role in security, as computing moves towards being evermore decentralized and distributed. The offloading function of a load balancer defends an organization against distributed denial of service attacks. It does that by shifting attack traffic from one area of the network to the other.

We implement 3 load balancing algorithms:

Least Connection Method:

It directs traffic to the compute component with the fewest active connections.

It is helpful when many persistent connections are not adequately distributed over components.

Least Response Time Method:

It directs traffic to the server with the fewest active connections and the lowest response time.

Round Robin Method:

It rotates servers by directing traffic to the first available server and then moves that server to the bottom of the queue. It is useful when servers have the same configuration and we do not have persistent connections.

4.3. Humans Storage Components

The human storage components are software-defined block devices that emulate the behaviour of a traditional block device such as a physical hard drive or an SSD. Generally, it is a form of network-attached storage.

Storage in the ecosystem is organized as blocks. It emulated the same type of behaviour seen in traditional disks or tape storage through storage virtualization.

Blocks are identified by an arbitrarily assigned identifier by which they may be stored and retrieved.

A filesystem is then applied on top of the block storage to “map” the files into a sequence of blocks.

4.3.1. Database File System

The database file system is used to store files closely associated with certain oracles or certain AI network engines or models, including images, videos, structured and unstructured data, documents, BL0B and CL0B files, etc.

The database file system provides significantly better security, availability, robustness, and scalability than traditional file systems. When data is stored in the database file system, they are backed up and synchronized with the blockchain. They are recovered along with their structural data in case of system failure. This makes storing data in the database file system appealing for many Dapps inside the ecosystem.

(image)

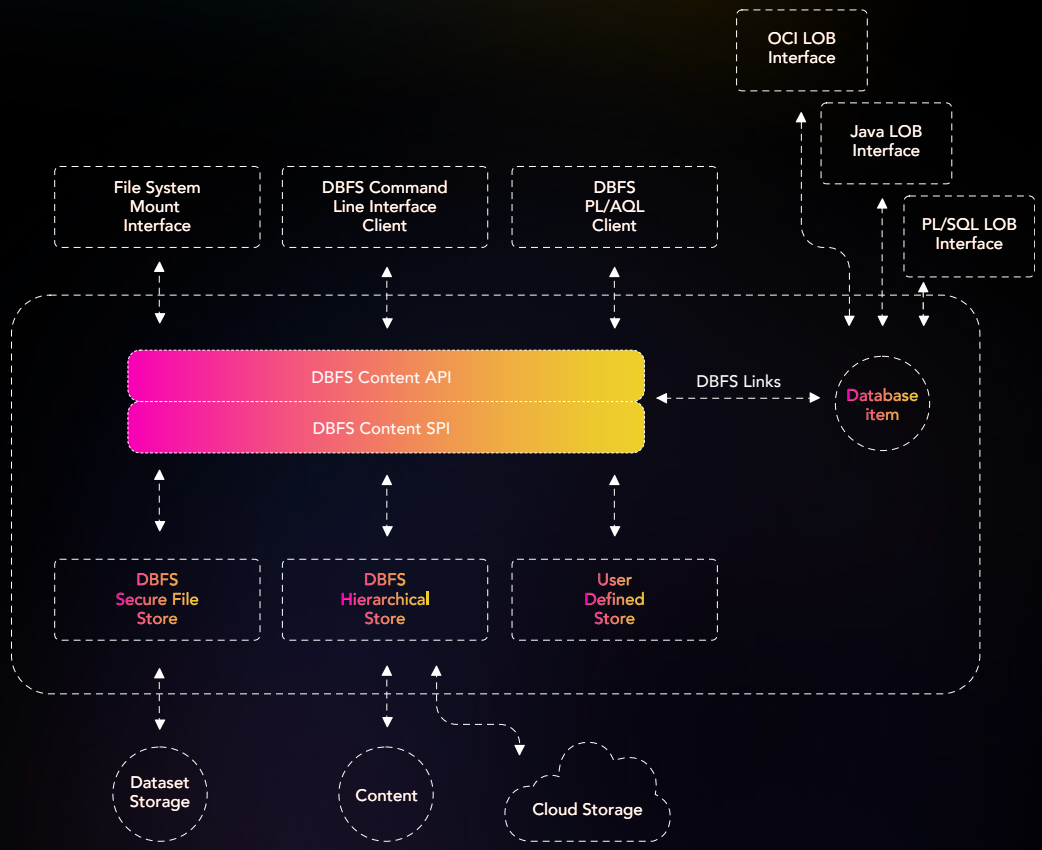

4.3.2. AI Network File System

Traditional file systems present significant challenges for artificial intelligence, machine learning and deep learning workloads.

Control node bottlenecks, inferior or non-volatile memory express drivers, and inefficient capacity utilisation are among the pain points that come with trying to process the millions or billions of small files that typically comprise AI, ML (machine learning), and DL (deep learning) workloads with legacy storage approaches.

The AI Network file system is built on top of the storage components engine to accelerate the data transfer between host systems and storage media, for example, over the PCI Express devices (peripheral component interconnect express bus).

It can enable our ecosystem to more fully utilize the storage maximum performance levels by increasing command counts and queue depth which is the primary concern for serving AI, ML, DL workloads.

4.3.3. AI Model File System

The AI model storage file system's architecture manages data as objects as opposed to other storage architectures managing data as a file hierarchy.

Each AI model is an object that typically includes the data itself, a variable amount of metadata, and a globally unique identifier. The model storage file system allows us the retention of massive amounts of unstructured data.

One of the design principles of the model storage file system is the abstraction of some lower layers of the HCC storage away from the other Dapps. Thus data is exposed and managed as objects instead of files or blocks.

These objects contain additional descriptive properties that can be used for better indexing or management. It allows us to address and identify individual models by far more than just names. The AI model storage file system adds a unique identifier across the entire system to support much larger namespaces and eliminate collisions.

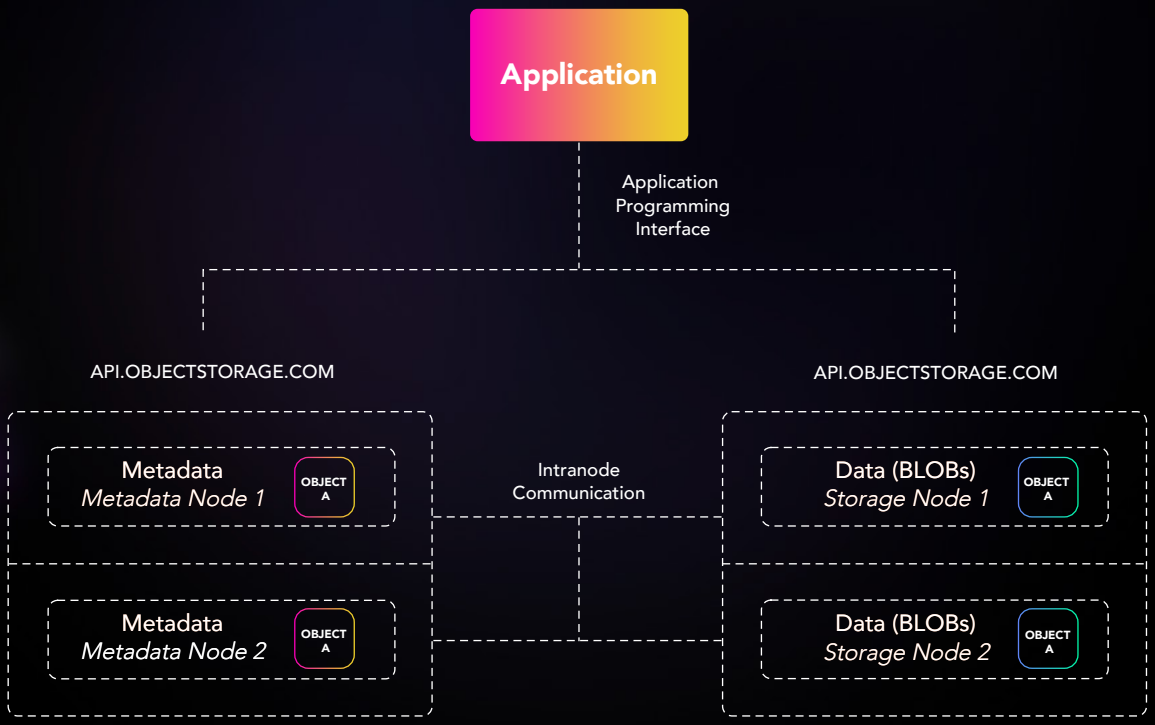

4.3.4. Blockchain File System

The blockchain file system is the underlying filesystem that stores the data of a blockchain compute engine. There are four pieces of data that are stored and maintained:

hum/blk*.dat:

The actual human blockchain blocks in the network format dumped in a raw format on the filesystem. They are used for rescanning missing transactions in a wallet, reorganising to a different part of the blockchain, or serving data to other compute engines that are synchronising.

---

hum/chainstate/*:

It is in-memory storage with a compact representation of all currently unspent transaction outputs and some metadata about the transactions they originate from. The data here is necessary for validating new incoming blocks and transactions.

It can theoretically be rebuilt from the block data; without it, in theory, one could still do validation, but it would mean a full scan through the blocks.

---

hum/index/*:

It is in-memory storage that contains metadata about blocks and where to find them on the filesystem.

---

hum/rev*.dat:

It is in-memory storage that contains “undo” data. They are blocks acting as “patches” to the chain state (they consume some unspent outputs, and produce new ones). They are necessary for rolling back the chain state, which is necessary in case of reorganizations.

4.4. Humans Compute Components

Compute components in the ecosystem are microservices that are encapsulated in containers (another type of virtualization).

In contrast to virtual machines, containers share operating system components, such as libraries and binaries, with the host operating system, whereas virtual machines do not. Thus, it allows a container to run thinner and with fewer hardware resources borrowed from the host. It offers a significant advantage in terms of speed and efficiency.

The ecosystem uses the CRI standard (Container Runtime Interface) to package a compute component and its dependencies in a virtual container that can run on any Linux, Windows, or macOS operating system.

It enables the compute component to run in various locations, such as premises, private or public clouds, by taking advantage of the resource isolation features of the operating system.

In the following diagram, we compare virtual machines and containers.

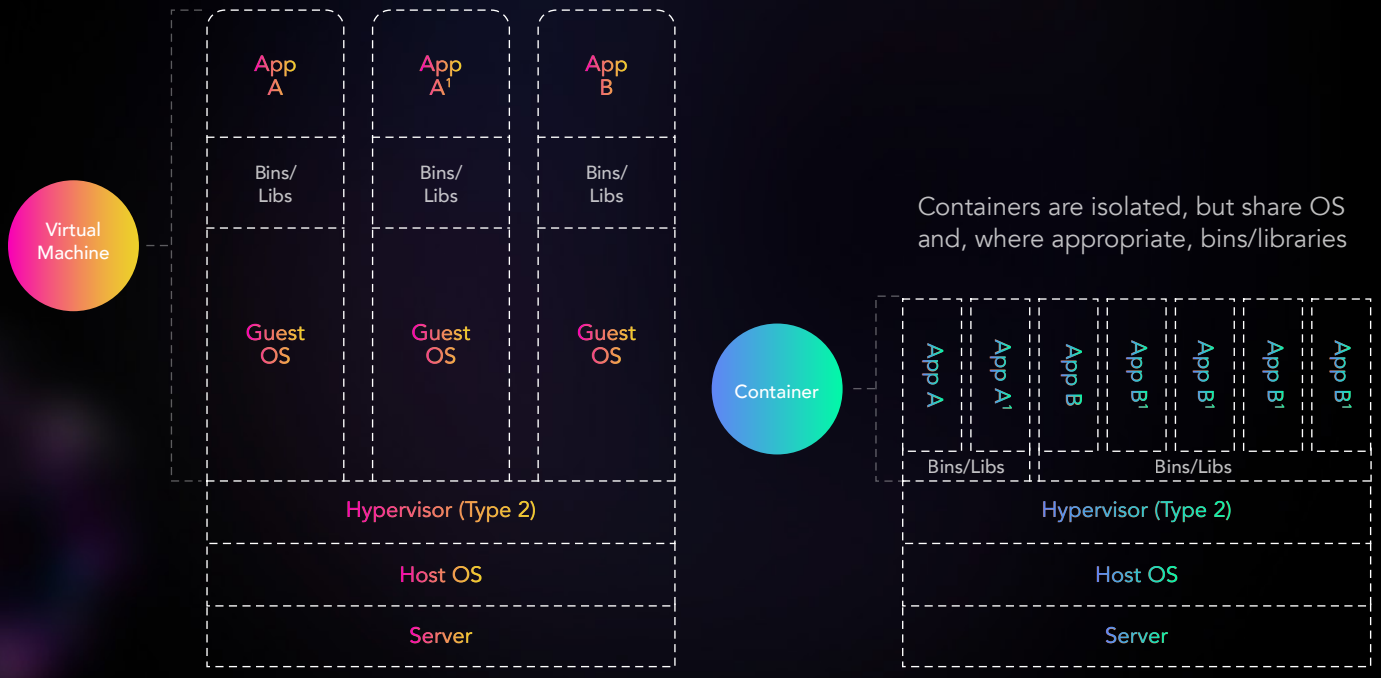

4.4.1. General-purpose compute engine

The general-purpose compute engine is a system in the ecosystem that allows us to deploy, scale and automatically manage compute components.

It is designed to run and scale millions of compute engines and deliver them consistently and efficiently irrespective of the complexity of the compute component.
It allows managing workloads on- and off-premises in hybrid or public clouds infrastructure, permitting the effortless movement of compute components.

The compute engine progressively rolls out compute components in the network while monitoring their health to ensure it does not stop all the instances simultaneously.

If a new compute component is going to be deployed at a future time in the ecosystem, it will be assigned its IP addresses and DNS names. Then, using a load balancer, it will direct the required traffic or requests between them without modifying an existing compute component.

If a certain compute node fails or dies, the compute engine will replace and reschedule the compute components, stopping those that do not respond to health checks. Furthermore, it will not advertise them to clients until they are ready to be served. One other advantage is that it places compute components based on resource requirements and other constraints without tampering with availability to mix best-effort and critical workloads to drive up utilization and save resources.

4.4.2. Encapsulation compute engine

The encapsulation compute engine is a plugin interface that allows the general-purpose compute engine that runs on every compute node in the network to use more than one specification of a compute component.

The encapsulation compute engine format contains sufficient information to launch the compute engine component on the target platform (e.g., commands, environment variables, arguments, libraries, etc.). In addition, this specification defines how to create a compute engine image, which is done by a specific build system that outputs metadata regarding a compute image manifest, a compute image filesystem, and a compute image configuration.

This specification also includes the content addressable identity of one or more filesystem serialization archives that will be unpacked to make the final operable filesystem.

The compute engine image also contains the configuration of the application arguments' environment, 3rd party dependencies, etc.

This combination of the compute image manifest, compute image configuration, and one or more compute image filesystem serializations is called the OCI image.

The compute image manifest has three main goals:

Create content-addressable images by creating an image model where the image configuration is hashed, generating a unique ID for the image and its components.

Allow compute images created using "big manifests" to add platform-specific prerequisites

(eg., Linux libraries, Windows registries, MacOS KEXTs).

To be translatable to any implementation of a 3rd party using this standard.

The compute image layer filesystem is a serialization protocol that supports serializing a filesystem and its changes like deleted files, created files etc. into a signed blob called a layer.

One or more blobs are applied on top of each other to create a complete filesystem.

The compute image is a collection of root filesystem changes in a specific order that have corresponding execution parameters for use within a compute engine runtime.

These specifications use a particular format JSON/YAML/PROTOBUF to describe images to use within a compute engine runtime and the execution tool and its relationship to the filesystem change sets described in layers.

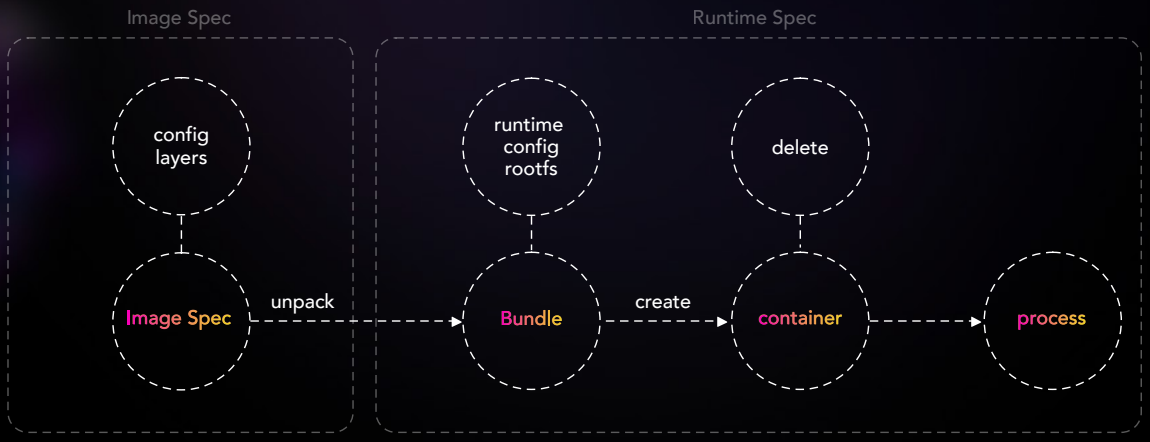

4.4.3. AI Network training engine

The AI network training engine is a compute component that encapsulates the machine-learning process at the level of a compute image to be executed by the general compute engine.

The process of the AI network training involves three steps: actual training, validating and testing. Each step requires a new dataset.

By feeding the initial data with tags and labels (metadata: notes and marks) into the AI network, the AI is trained to produce a particular prediction with each run. Each time, the parameters can be adjusted to ensure that the predictions become more accurate with each training step.

The algorithm is then validated by running another dataset with tags and labels against the trained model. New variables may need to be adjusted to improve the algorithm at this point by comparing the output validation data with the training data.

Once it passes the validation stage, the system can be tested with new data with no tags or labels, confirming if the network is ready to be used for its intended purpose. For example, if we train the algorithm to produce voice reading a text, we would feed it datasets of texts and recordings of voiceovers. And for the test stage, we would feed it texts only to examine the voice results produced by the algorithm.

After the confirmation is received, the network is then signed into a general compute image and deployed using the general compute engine onto the network, thus making it usable to the whole network for creating new models.

Some oracles might use this network to make it available to the marketplace or any other third-party system.

4.4.4. AI Model generation engine

An AI model is a file that has been trained to recognize certain types of patterns. We train a model over a specific set of data, e.g., someone's voice or footage of a person, providing it with an algorithm that it can use to learn from this data.

Once we have the model trained, we can use it to infer results based on data that it has never seen before.

In summary, the model represents what was learned using the AI network on our training data and represents the rules, numbers and other algorithm-specific data structures required to make predictions.

4.4.5. Blockchain engine

The term “blockchain engine” is being used in the ecosystem in direct relation to the term “blockchain”, a decentralized digital ledger that records all cryptocurrency transactions and makes the information available to everyone via a connected device. Every transaction will be chronologically recorded and distributed to another pool of “blockchain engines”. These “blockchain engines” communicate with each other and transfer information about transactions and new blocks.

Blockchain engines are a critical component of the blockchain infrastructure. They maintain the security and integrity of the network. These blockchain engines have the primary purpose of verifying each batch of network transactions, also called a block. They are distinguishable from other engines by a unique identifier.

The ecosystem has two types of engines: full engines and light (listening) engines. The full engines are the primary support for the security and integrity of the network. They download the blockchain’s entire history to monitor and enforce its rules.

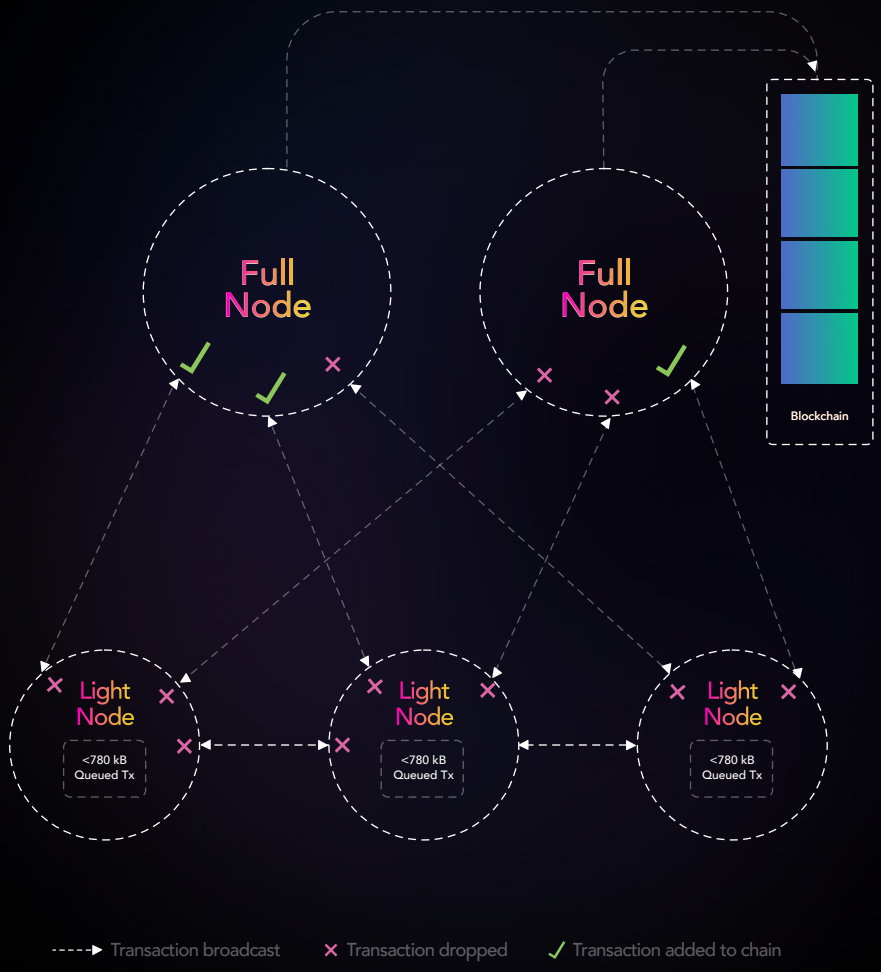

Each connected user or party in the ecosystem is a light blockchain engine. The light blockchain engine has to connect to a full blockchain engine to participate. It is essentially a public proxy to a full blockchain engine. It connects with any other blockchain engine that agrees to have a connection. A reliable full node usually runs 24/7, relaying information and blockchain history to multiple other nodes.

4.4.6. Blockchain oracle engine

Blockchain oracle engines are services that provide smart contracts with external information and functionality. They serve as bridges between multiple blockchains or external services.

Blockchains and smart contracts cannot access off-chain data usually. However, it is occasionally mandatory to have relevant information off-chain to execute a specific agreement or condition. Here the blockchain oracle engine comes into play, as they provide the ecosystem with a link between off-chain and on-chain functionality.

These engines play a vital role in the ecosystem because they broaden the scope in which our smart contract engine can operate. Without these blockchain oracle engines, the contracts will have limited use as they would only have access from and within the blockchain.

One important thing to mention is that the blockchain oracle engine is not a data source but rather a layer that queries, verifies and authenticates external data sources and then relays that information. The data transmitted by oracles can be in many forms request/response, action, action completion, etc.

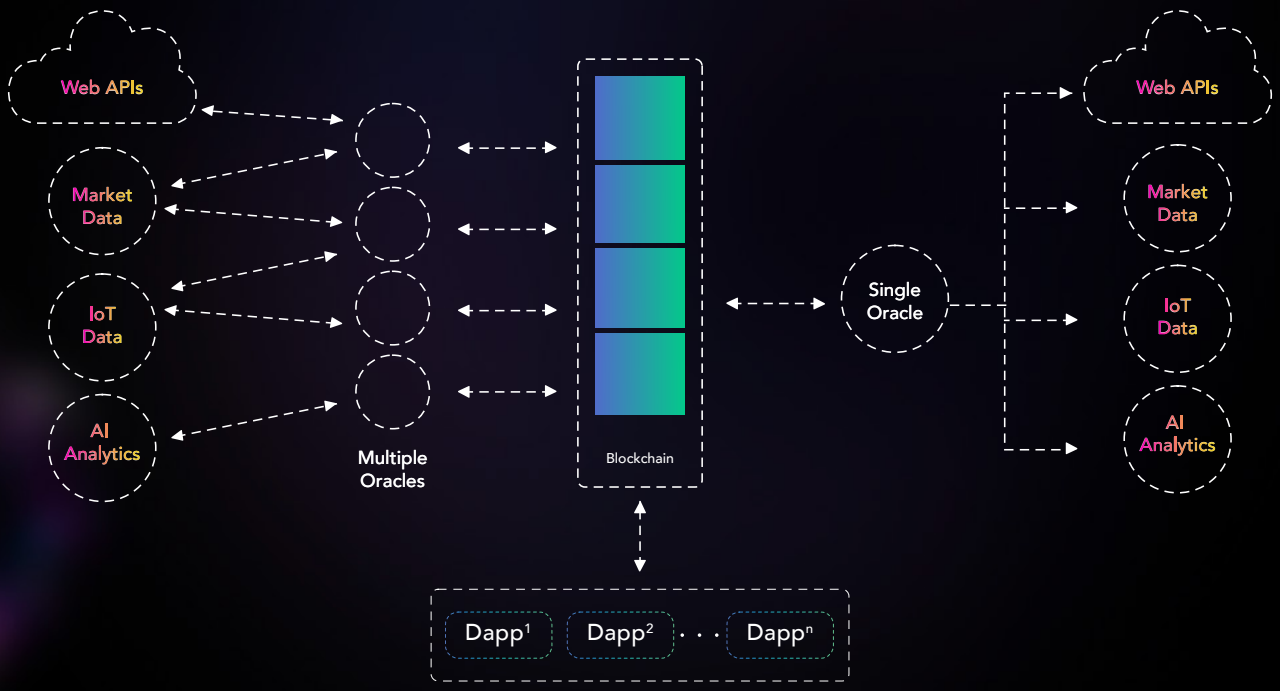

4.5. Humans Smart Contracts

4.5.1. Hybrid NFTs

The purpose behind developing our own double token standard is to accommodate the minting, managing and governing of the AI NFTs, at scale.

4.5.1.1. Overview

AI NFTs represent a new asset class with the sole purpose to encapsulate neural networks and their variations (trained neural networks with user-generated data) and provide a command type interface on top of it.

Essentially the design of the standard takes into consideration the following components:

The asset has a non-fungible part representing the link to the hashed neural network.

The asset also has a fungible part representing different rights in managing and potentially governing the AI NFT:

Each AI NFT can be opened as DAO by the minter.

Permissionless access to the AI NFT is ensured through the DAO-like structure:

Participants own the AI DAO NFT's fungible part.

4.5.1.2. Components

The new asset class we are creating will serve several functions. Based on the functions to be served, we will split the components into two main categories:

1. Non-fungible

2. Fungible

The non-fungible part of the asset behaves like an ERC-721 and plays the following roles:

- Ownership
- Link to the original neural network
- Transfer of ownership

The fungible part of the asset works like an ERC-20 token and plays the following roles:

- An entry point for HV block producers to validate transactions and earn a split of the profit based on their minted blocks
- Governance in the AI NFT DAO

4.5.1.3. Validation

HV Block Producers stake the fungible part for the right to validate requests. The validation is performed as part of the BFT-based consensus on which the block producers will reach consensus if to sign a block or not.

Ideally, the request processing will be split into two, separating the validation from the compute and as such performing validation in batches independently.

4.5.1.4. Stake

The stake plays a key role in having every entity involved committed and aligned financially.

Data providers need the fungible part of the asset:

To encompass the governance of the AI NFT:

The community may be required to lock the tokens for governance rights.

The validation part is crucial to manage the asset:

The HV block producers are required to stake the token.

The HV will earn a portion of the revenue for which they perform their tasks.

4.5.1.5. Authorization

Human Validators play a key role in authorizing requests received by the AI NFT. The data provider will rely on the community to enforce the settings decided upon.

The AI NFT owner can intervene if necessary and modify the approval rules, the minimum stake required, etc., and adjust the financial incentives for HV.

4.5.1.6. Governance

In governance, the goal is to reach a consensus on the settings that Human Validators must enforce. The settings are decided as follows:

AI NFT:

by the data provider;

AI DAO NFT:

there are cases when the data product creator opens governance to the community;

the fungible part of the asset class represents the governance rights;

formal on-chain governance is both entity-based (based on identity confirmation) and stake-based (based on the token stake).

4.5.2. Contracts architecture

Proof-of-Human is the first widespread application for AI using zk-SNARKs, a novel form of zero-knowledge cryptography.

The acronym zk-SNARK stands for "Zero-Knowledge Succinct Non-Interactive Argument of Knowledge" and refers to a proof construction where one can prove possession of certain information, e.g., a secret key, without revealing that information and without any interaction between the prover and verifier.

"Zero-knowledge" proofs allow us to prove that several statements of our architecture are true without revealing any information beyond the validity of the statement itself.

In a zero-knowledge "Proof of Knowledge", the requester can convince the verifier not only that the request exists but that they, in fact, know such a request – again, without revealing any information about the request.

Zero-knowledge proofs can be verified within a few milliseconds, with proof of length of up to 1K bytes even for statements about requests that are very large. The most efficient known way to produce zero-knowledge proofs that are non-interactive and short enough to publish to a blockchain is to have an initial setup phase that generates a common reference string shared between requester and verifiers. This reference is called the public parameter of the request.

4.5.2.1. zk–SNARKs entities

To have zero-knowledge proofs in Humans, the function or smart contract that needs to determine the validity of a request according to the request/smart-contract rules must interpret whether the request is valid or not, without revealing any information on how the calculations are performed. This is done by encoding some of the verification rules at a machine execution level by initially converting the equivalent of knowing a solution to what needs to be proved into an algebraic equation.

The Humans ecosystem applies a simpler zero-knowledge proof, following the logic presented below:

A digital signature on the data provided by a user towards an AI Compute Model can be thought of as a very specific, narrow example of a zero-knowledge proof.

The user provides a message M and a digital signature over M (hash of M encrypted with the user's private key) to the oracle.

The Oracle can hash M for themselves, decrypt the provided signature with the user's public key, and compare the two. If they match - there is a confirmation of the signature. Thus, the oracle verified the arithmetic "proof" that the user provided over the following NP (nondeterministic polynomial) time statements:

Oracle knows the user's public key corresponding with the private key used by the user to sign the message M.

This way, it is proven that the user knows the private key. Some may argue that signatures, in this case, are not an example of zero-knowledge proof of knowing the private key because the oracle has learned more than just the fact that the user knows the private key. The oracle can convince another oracle that he knows the user's private key just by passing the user's signature.

However, the zk is the user's private key in this case. The user never reveals the private key to the oracle, but the oracle is convinced that the user must know the key because only with it, the user could produce a set of all three: the message, the public key, and the signed hash.

Given enough iterations and cases (valid and invalid signatures), anyone could verify if the message M was indeed signed with the zk or, in other words, the user's private key.

Even though zk's do not necessarily guarantee non-transferability, which is a different property, with the appropriate techniques, we can make the signature non-transferable.

With this idea of the zk in mind, we can generalize it to our use case in the Humans ecosystem to a computable NP statement expressed logically as an arithmetic circuit.

For example: User signs data (e.g. input data for AI model) => Oracle verifies data using the signature => AI model validator signs data => Oracle verifies data => Oracle triggers AI model execution => AI compute module verifies the User's data and the validator's data => AI compute module produces output => AI compute module signs data => User verifies output data.

Below we give a brief overview of the entities involved, their role, and the rules for determining a valid request to be transformed into an equation that then can be evaluated without revealing any sensitive information to the parties verifying the request.

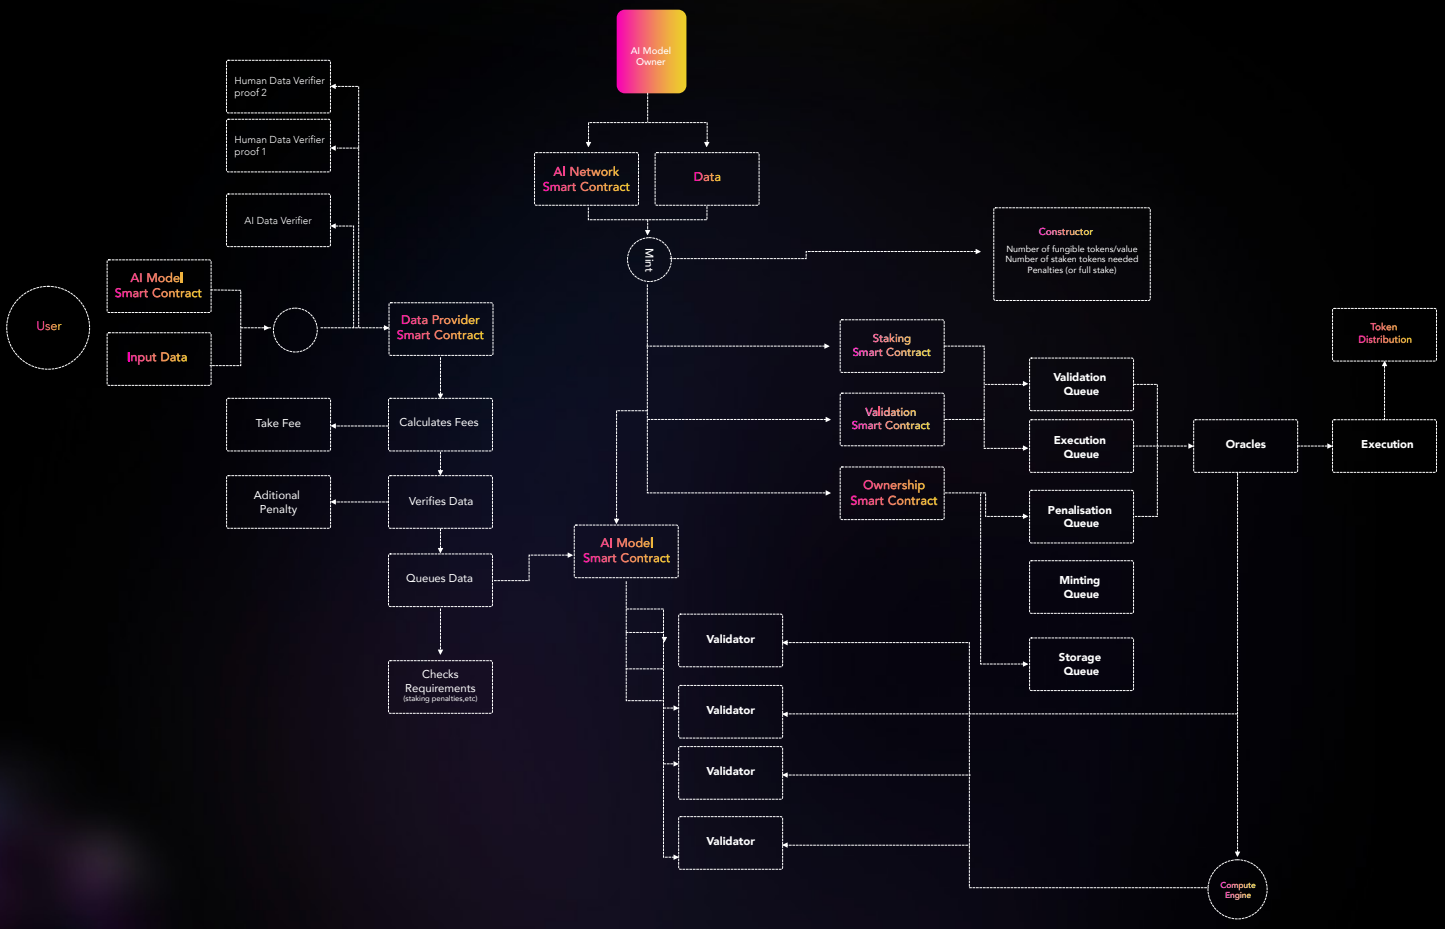

(image)

4.5.2.2. AI Network Contract

AI Network Contracts learn (or are trained) by processing examples, each of which contains a known "input" and a "result", forming probability-weighted associations between the two, which are stored within the data structure of the AI network itself.

Such AI Network Contracts "learn" to perform tasks by considering specific input data, generally without being reprogrammed using input-specific rules. For example, in image recognition, they might learn to identify images that contain cats by analyzing input data of images that contain hotdogs, providing a result of "hotdog" or "not hotdog" (pun intended). They do this without prior knowledge of cats. We use this as a smart contract to provide a unique ecosystem where researchers and developers can create neural networks, and different actors can use these networks to create unique AI Model Contracts.

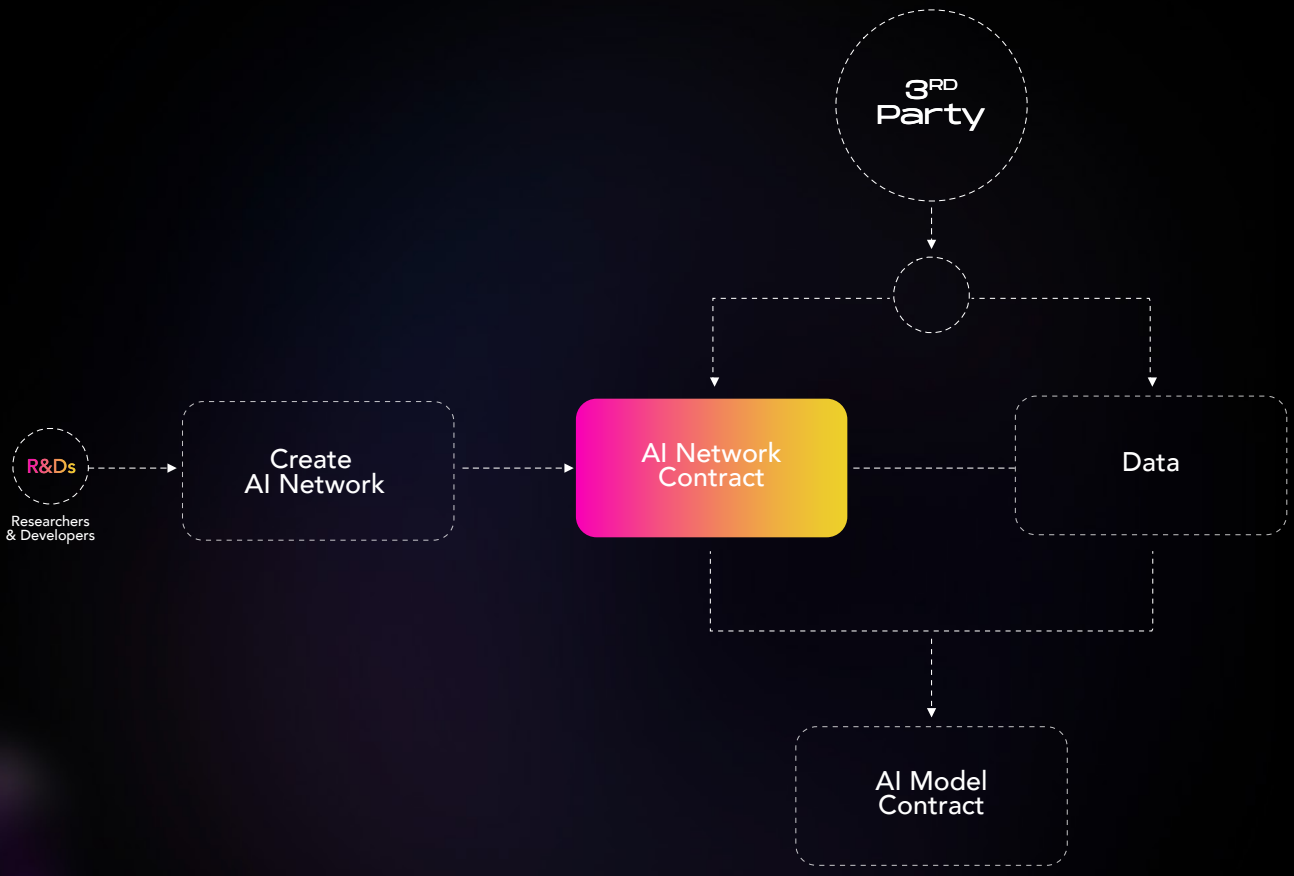

4.5.2.3. AI Model Contract

The product of an AI Network Contract and data from a training set that can be as varied as a corpus of text, a collection of images or sounds, data collected from an individual or a service, is called an AI Model.

Such AI Model Contracts are as unique as the provided input data to the AI Network Contract. These models derive from the AI Network contract by using unique input data, such as someone's voice or video footage of someone's face.

It allows us to use the power of the AI and the uniqueness of each training data to provide unique services like "voice generation", "spoke-person generation" of a specific 3rd party without their intervention.

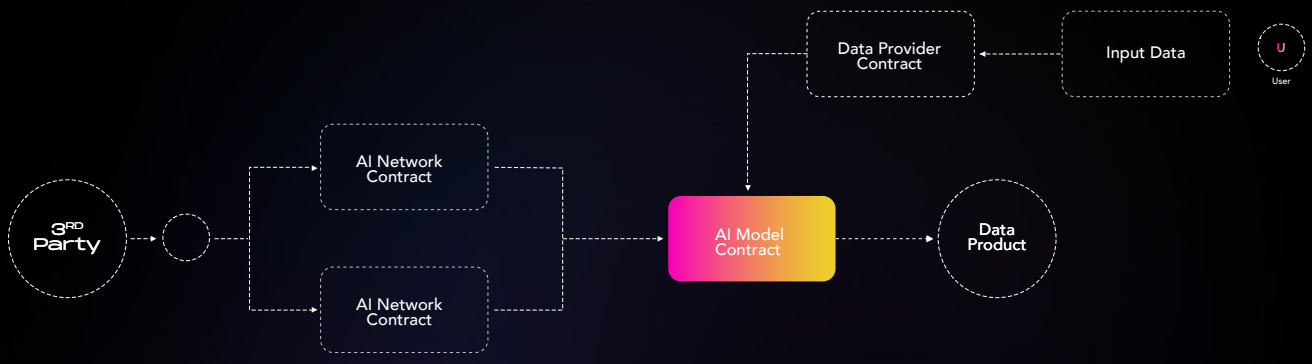

eg.,: Morgan Freeman provides recordings of his voice to a voice AI Network Contract generating an AI Model Contract;

The client provides a text to the AI Model Contract resulting in a Data Product with a recording of Morgan Freeman's voice reading the text.

4.5.2.4. Data Provider Contract

To assure the ethical integrity and security of the contract, we offer a new Humans Improvement Proposal that filters data sent to the AI Model Contract.

The Data Provider Contract communicates with the Staking Contract to verify that the staking requirements of the AI Model Contract are satisfied. It also communicates with the Validation Contract to queue an additional human verification to the validation parties specified in the AI Model Contract.

The Data Provider Contract also has a posterity role for contracts as their request data and decisions will also be recorded in this contract.

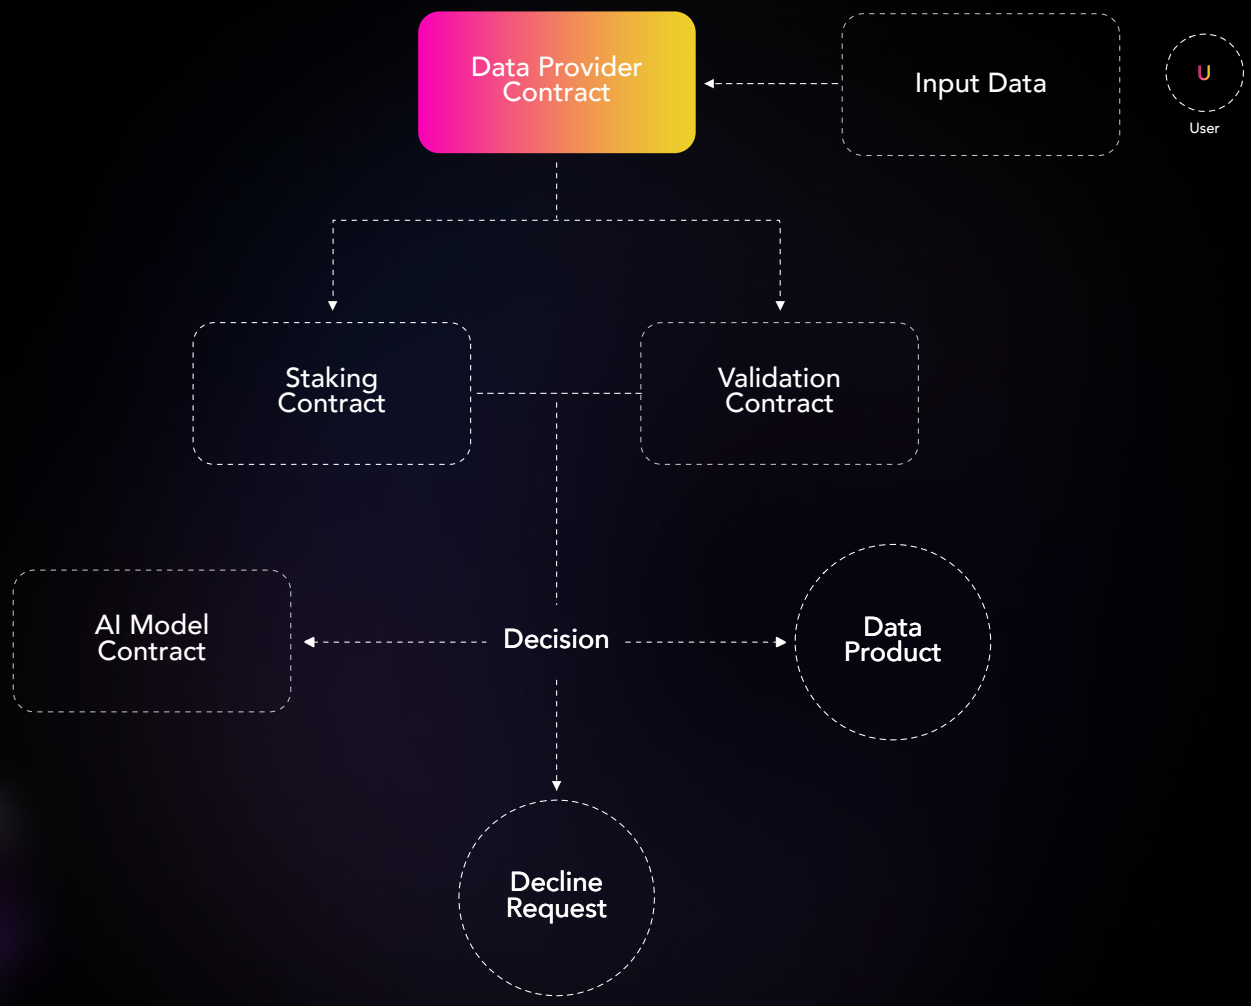

4.5.2.5. Staking Contract

Validators and users of the contract may need a stake to use certain AI Model Contracts.

By requiring a stake from the validators, the contract assures that they act in accordance with the rules imposed by the contract owners, disincentivizing foul play through financial penalties.

As the validators must respect the contract rules, users of the smart contract (e.g., third party apps, other contracts) must too respect the rules of the AI Model Contract for any external usage. So by requiring a stake, we assure that they also play fair.

The Staking Contract communicates or receives decisions from the Validation Contract, AI Model Contract or certain oracles that will have an input on what happens with the stake of the user.

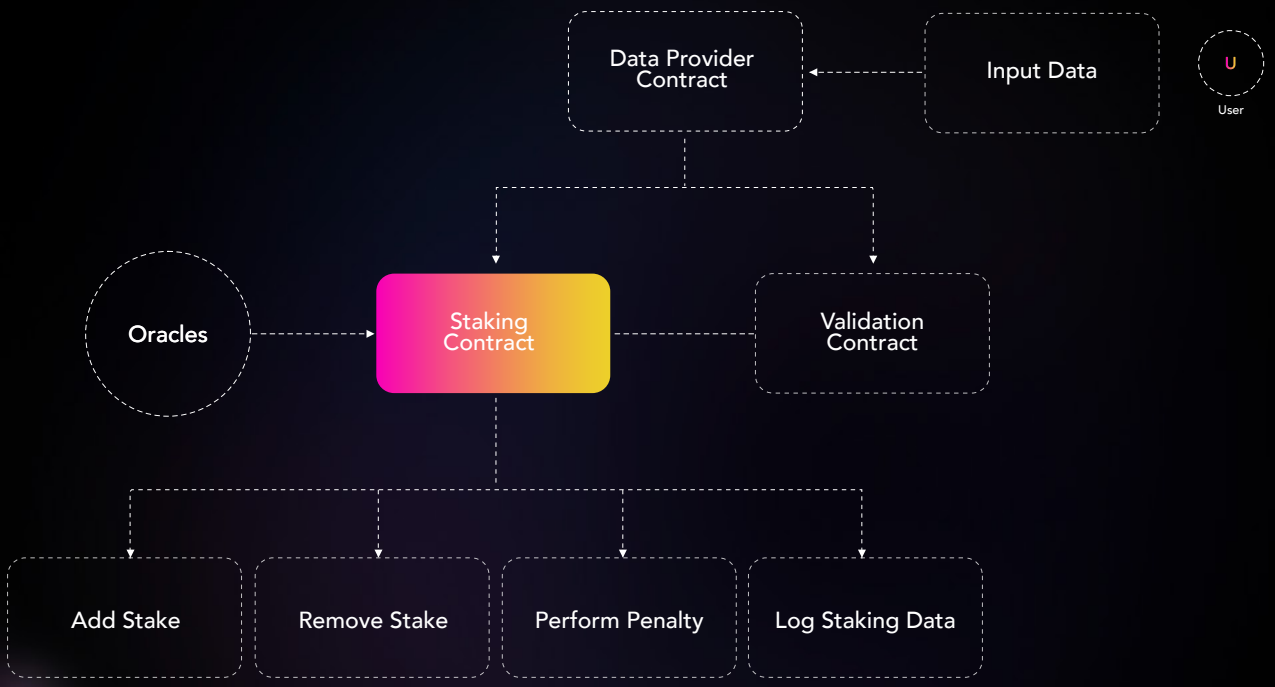

4.5.2.6. Validation Contract

The Validation Contract is a filtering layer between the request of a party and an AI Model. The contract also has the role of proxying requests toward oracles that validate request data automatically and recording all these requests and decisions for posterity.

It also communicates with human validators of the AI Model Contract to get an additional level of validation if the contract specifies it.

When a decision is made, the Data Provider Contract notifies relevant oracles to inform the AI compute engine that the computation of the Data Product can begin.

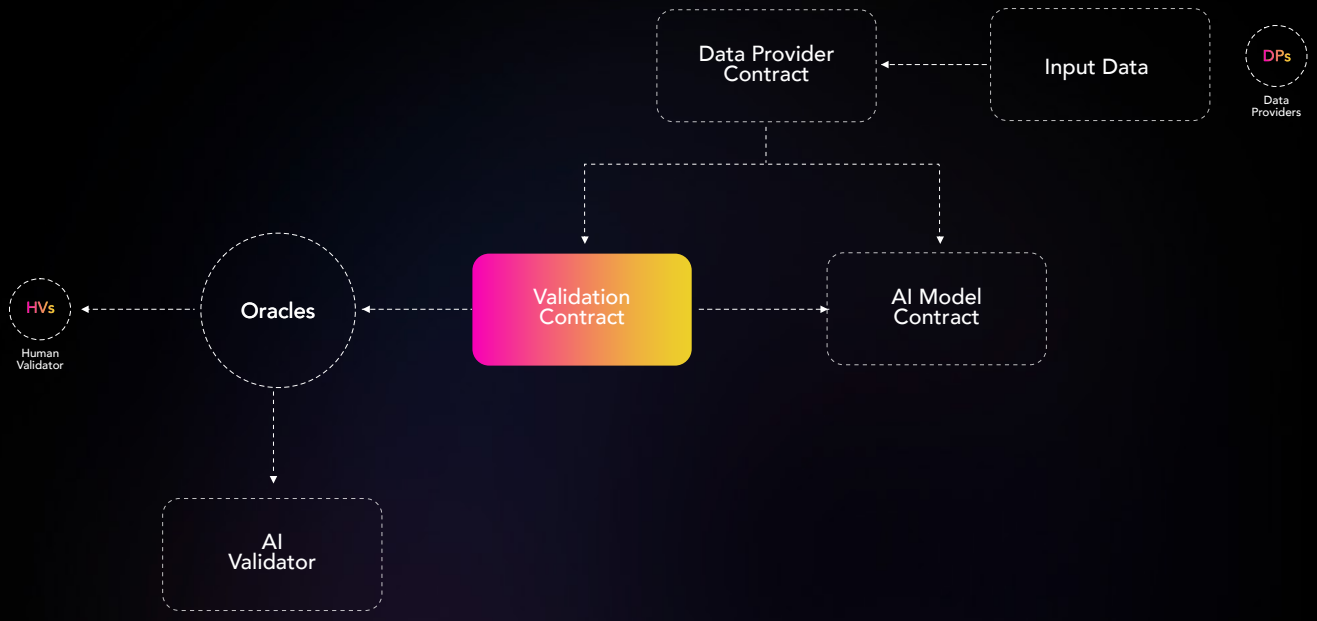

4.5.2.7. Ownership Contract

The Ownership Contract records all the historical events of a contract for posterity. It has the role of a ledger where any change to the ownership of a contract is announced on the network for other contracts or oracles to find the new ownership of the said contract. Generally, the Ownership Contract serves as a sidechain on the blockchain, where those who have a contract can log in. It can also retrieve all the contracts of a specific owner or, alternatively, all the owners of a contract.

New functionality may be added further down the road, such as multiple owners, ownership voting, ownership transference, ownership stake, etc.

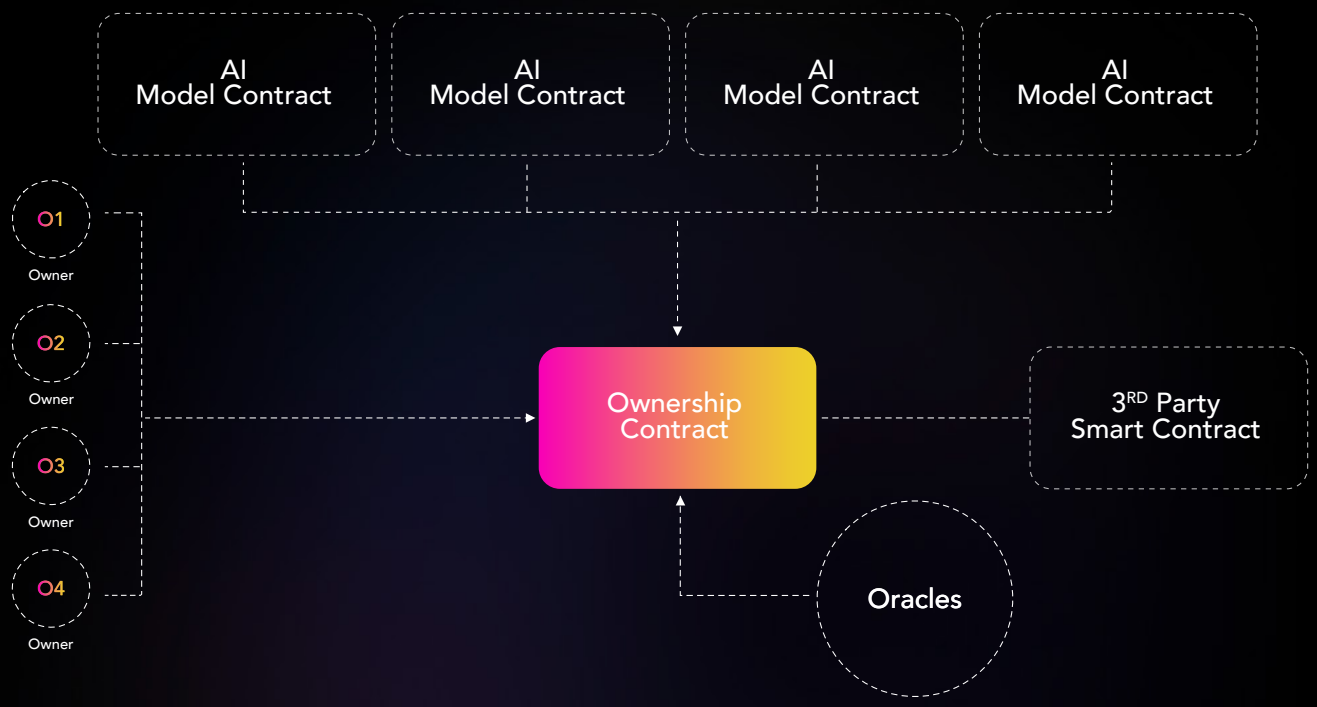

4.5.3. Humans ecosystem oracles

A Humans oracle is a system inside the Humans ecosystem that connects smart contracts with different systems to feed information in the blockchain and reverse.

Information from the ecosystem encapsulates multiple services/sources so that decentralized knowledge is obtained. Information to the ecosystem may include “ai network execution”, “ai network training”, “ai model execution”, “ai network storage”, “ai model storage”, “human validator notification”, “ai validation model execution”, “notification services”, etc.

The oracle is the layer that queries, verifies and authenticates via zkSnarks data sources, trusted APIs and then relays that information to the blockchain.

Concern:

If an oracle relies on a single source of truth (centralized) that can lead to issues: the data source can be hacked in a MitM Attack (man in the middle) or altered by its owner to sway smart contracts. zkSnarks oracles increase the reliability and security of the information provided to smart contracts by querying multiple data or execution sources. It achieves near trustlessness since oracles are a part of the ecosystem and use zkSnarks of the contracts as a security mechanism.

4.5.3.1. zkSnarks construction in Humans ecosystem

In order to have zk (zero-knowledge) execution in Humans ecosystem, the function determining the validity of the request according to the Contract constraints must return the answer of whether the request is valid or not, without revealing additional information.

It is done by encoding some of the request metadata rules from the contract in zkSnarks. At a high-level view, zkSNARKS work by first precisely converting what we want to prove into a homologue form about knowing a solution to a mathematical equation. We will explain further a brief overview of how the rules for determining a valid request get transformed into an equation that can then be evaluated and verified without revealing sensitive information to the parties verifying the equation.

Computation -> Mathematical Circuit -> R1CS (Rank 1 Constraint System) -> zkSnark

It is an example of how a mathematical circuit looks for computing the function
*F = (A+B)xBxC ****

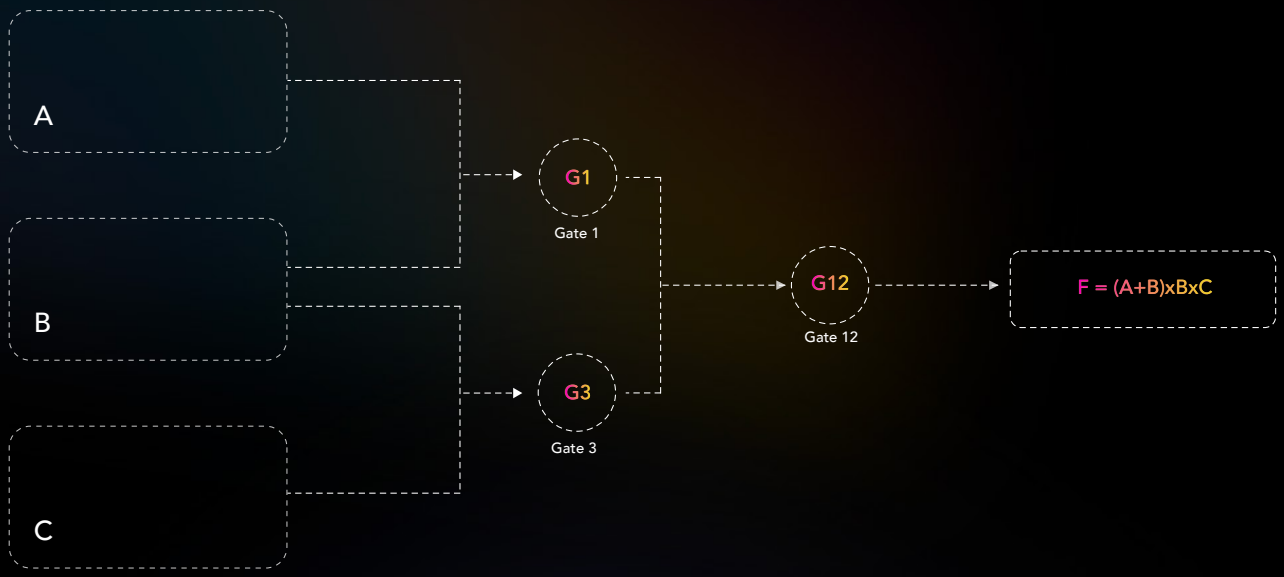

In this mathematical circuit, we observe that the input data (A, B, C) are propagated from left to right. In this case, the R1CS (Rank 1 Constraint System) confirms that the value from a certain gate is either multiplication or addition. In this example, Gate 1 is (+) addition, Gate 2 is (x) multiplication, and Gate 12 is (x) multiplication.

In this R1CS representation, the verifier has to check a constraint for each wire of the circuit. This method uses a representation of the circuit QAP (Quadratic Arithmetic Program).

It flips the constraint to be checked from numbers to polynomials. As polynomials can get quite large, the identity does not hold between them, and it will fail to hold at most points. This means that we have to check only two polynomials that match at one randomly chosen point to verify the proof with high probability.

Suppose the requesting party (prover) had prior knowledge of the point the verifier would check, they could craft polynomials that are invalid but satisfy the identity at that point. However, in zkSnarks, homomorphic encryption and elliptic curve pairings are used to evaluate polynomials in the dark without knowing the evaluated point. Public parameters determine which point to check, which is executed in an encrypted form, so neither the prover nor the verifier knows which point it is.

5

Ecosystem Economy

5. Ecosystem economy

5.1. Token Allocation

$HEART tokens offered for the public sale will be distributed as ERC20 tokens and further substituted by $HEART tokens on the native Humans blockchain after its launch in 2022.

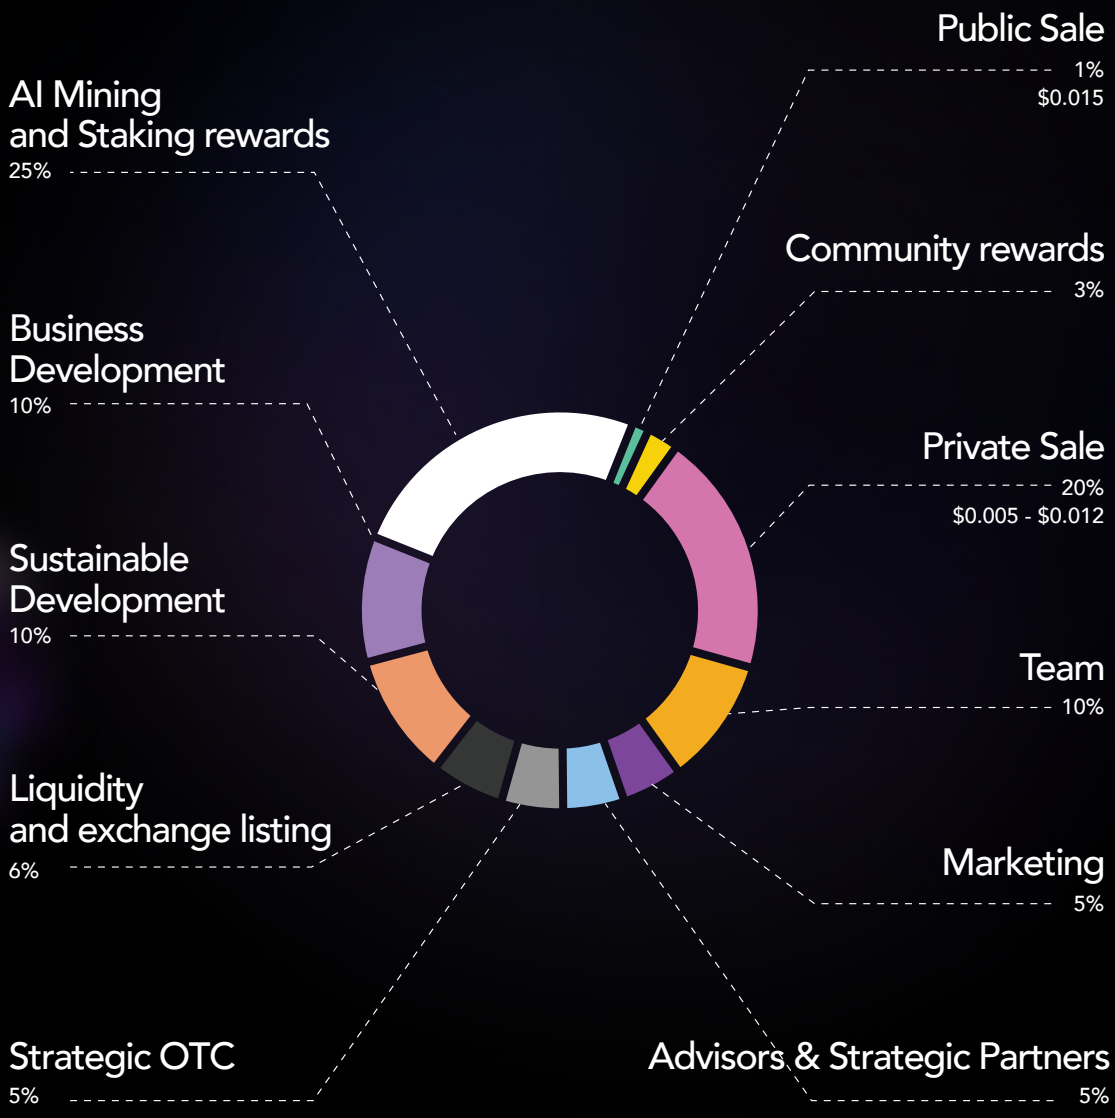

HUMANS $HEART
Unlock Schedule

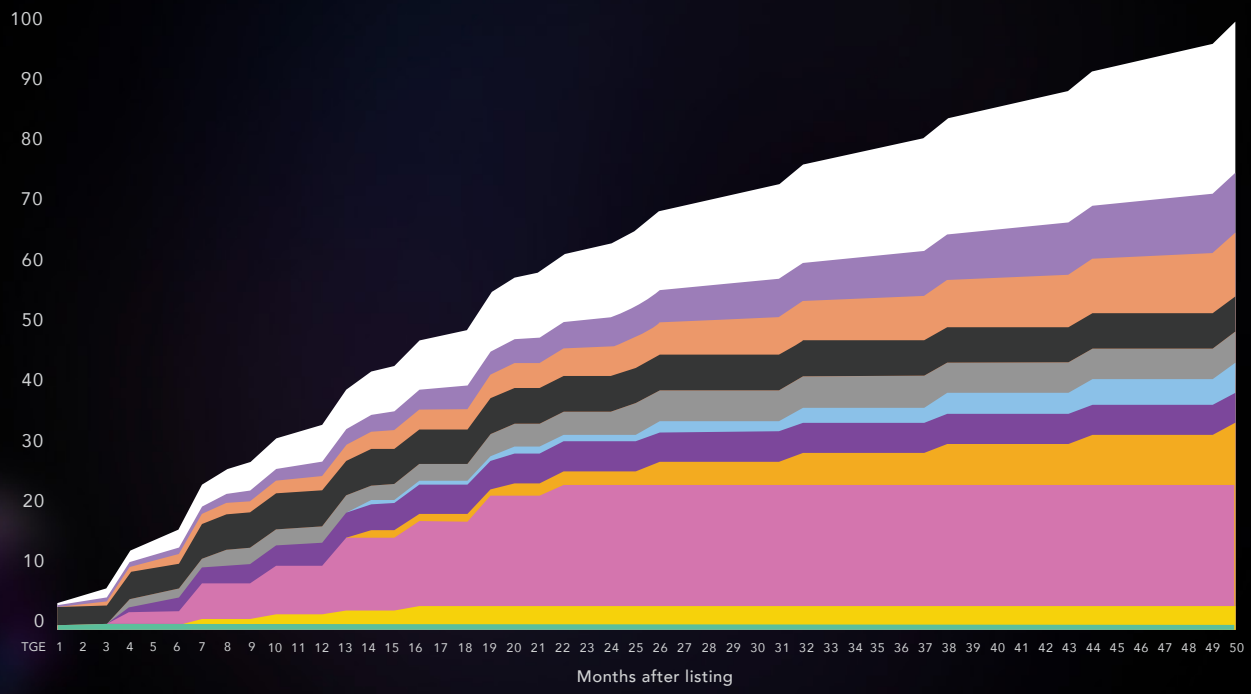

6

Roadmap

6. Roadmap

The Humans project was launched in 2020, and the most important piece of its technology, the Humans Studio, has already launched in its beta version. The first project showcasing the technology, Tovia.ai, is live and onboarding B2B clients. We are developing the rest of the ecosystem, the details of which are illustrated on the timeline below.

Our history is a story of creators.

Watch the movie ▸ https://youtu.be/qGgSJ_C23XM

7

About

7. About

7.1. Company history

Humans is a deep tech start-up creating a blockchain-based platform integrating data and technology to generate unique AI models and allow the business community to use them as an API (application programming interface).

The company was born on March 2, 2020, when we raised $400k from 2 venture capitals (VCs) and a media business investor. This money allowed us to validate the concept by selling synthetic media created on our platform that uses partners' AI trained on users' data to clients. Being laser-focused on integrating tech from research partners and our technology as opposed to researching from scratch, allowed us to create synthetic humans in a record time of just a few months.

Our initial focus was on building a technology platform to generate and manipulate the Digital human DNA, creating synthetic media. We integrate data and partner AI technologies in one platform to create Human Digital DNA and open it as an API to businesses. Our product is a technology platform able to generate and manipulate the human Digital DNA, creating synthetic media. Our vision is a platform in which anybody can create anything, driven by AI.

More specifically, based on the latest advancements in AI, the Humans' core technology consists in manipulating one's voice, appearance, movements, and gestures to generate synthetic representations of one's identity - or what we refer to as the digital DNA. With this digital DNA, the users of the Humans platform can synthetically generate any type of media: audiobooks narrated by actors, advertisements edited post-production, audio-video personalization in any language, etc.

In December 2020, we set out to work creating a community of researchers and academics around Humans by organizing the world's first synthetic media Summit. Thus, we opened the conversation about their needs (monetization, access to data, technology integration, cloud architecture, being part of a bigger product and scaling the human potential).

In June 2021, Humans launched a biometrics solution for its blockchain-built platform that uses AI to create digital DNAs. The Proof-of-Human tech innovation was announced at the Dubai AIBC Summit, a global event on emerging technologies.

In September 2021, Humans attracted USD 9M financing through the private sale of $HUMAN tokens. Among the main investors for this new round are the entrepreneur Răzvan Munteanu, one of the most active investors in the blockchain area, and Elrond Research, the investment arm of one of the most respected blockchain companies in the world.

Humans is backed up by a hyper-specialized team (40+ people), consisting of highly reputed professors, blockchain experts, software experts, computer vision PhD students and international Olympic students.

7.2 Team

Humans' team includes over 40 devoted members. 10% of $HEART tokens allocated to the team will be vested for 3 years.

Core Team + 30 more Superhumans

Sabin Dima
CEO

Razvan Munteanu
Head of Strategy

Nicu Sebe
Head of AI

Florin Otto
Head of Crypto

Lucian Nicolescu
Head of User Product

Vali Malinoiu
Head of Infrastructure & Blockchain

Dragos Bunescu
Head of Business

Dragos Costea
Head of Vision

Valentin Marcu
Head of Mobile

8

Definitions and references

8. Definitions and references

The terms used throughout the paper have the following definitions:

AI = Artificial Intelligence, the simulation of human intelligence processes by machines, especially computer systems.

AIs = within the paper mean results of training Algorithms with Genomes or other Data.

AI Model Contract = AI NFT = NFT representing a trained AI algorithm.

AI Developers / Researchers = the entities deploying algorithms to the AI Library so they may be used by the community based on the conditions they set.

AI Generation Oracle (HAGO) = an oracle starting the process of using a neural network to produce unique AI models based on specific training data.

AI Model File System (HAMFS) = data storage architecture that manages data as objects, as opposed to the other storage architectures, like file systems, which manage data as a file hierarchy.

AI Model Generation Engine (HAMGE) = engine that takes an AI network with some training data, resulting in a file that generates patterns based on that data.

AI Network File System (HANFS) = is built on top of the storage components engine to accelerate the transfer of data between host systems and storage media, over the devices peripheral component interconnect express bus.

AI Network Training Engine (HANTE) = a compute component that encapsulates the Machine learning process at the level of a compute image to be executed by the general compute engine.

Algorithms = programs developed to perform certain AI Personal profile and biometric data, e.g., functions, prior to training with data users' voice, likeness etc.

App Builder = module for building recipes from the data products that are available.

App Builders = users that take advantage of existing AI models or recipes in the platform and expose the information to customers in proprietary ways.

Apps = a combination of AIs that generate valuable outputs on request.

Biometrical data = personal data resulting from specific technical processing relating to the physical, physiological, or behavioural characteristics of a natural person, that allows or confirms that natural person's unique identification of, such as facial images or fingerprint data.

Blockchain engine (HBE) = is being used in the ecosystem in direct relation to the term blockchain, a decentralized digital ledger that records all cryptocurrency transactions and makes the information available to everyone via a connected device; every transaction will be chronologically recorded and distributed to another pool of "blockchain engines",

these “blockchain engines” communicate with each other and transfer information about transactions and new blocks.

Blockchain File System (HBFS) = the underlying filesystem that stores the data of a blockchain compute engine.

Blockchain oracle engine (HBOE) = services that provide smart contracts with external information and functionality, serving as bridges between multiple blockchains or outside services.

Compute Nodes = Virtual machines running distributively across many computers that handle all kinds of compute requests, from oracles to training AI Models, and running AI Networks.

Consensus = A consensus mechanism is a fault-tolerant mechanism that is used in computer and blockchain systems to achieve the necessary agreement on a single data value or a single state of the network among distributed processes or multi-agent systems, such as with cryptocurrencies.

Container Network Interface = a virtual network component that acts as a network adapter between compute engines.

Humans Studio = the entry point of data providers, users, and app builders in the Ecosystem, wherein they can choose an algorithm, upload their data, and train their algorithm.

Data = virtually any kind of data set.

Data Product Generation Oracle (HDPGO) = a system that triggers the start of a data product generation process based on certain events (e.g., data validation through proof-of-human, different calls from other oracles, third-party integrations, etc.)

Data providers = users that provide raw data to train and customize existing AI algorithms.

Data Validation Oracle (HDVO) = an asynchronous oracle that provides other compute oracles or engines with binary decisions regarding input data of any kind.

Database File System (HDFS) = used to store files closely associated with certain oracles or certain AI network engines or models, including images, videos, structured and unstructured data, documents, invoices maybe, BL0B and CL0B files, etc.

Deepfake = a video of a person whose face or body has been digitally altered to appear to be someone else, typically used maliciously or to spread false information.

Digital genome / Digital DNA = all the data points defining how a human looks, sounds, and acts, capable of generating synthetic digital representations of his identity, known as synthetic or algorithmically created media.

Egress Network Engine = in our ecosystem implies traffic that exists on an entity or a network boundary, while the Ingress network engine enters the network boundary.

Encapsulation compute engine (HECE) = a isolation / jail process where certain compute engine processes can run using their own manifests (e.g., operating system, libraries, dependencies), without tampering with other compute engine processes executing on the same node.

Governance tokens = the rights to govern the NFT and the rights to any $HEART revenues it may generate.

General-purpose compute engine (HGPCE) = a system in the ecosystem that allows us to deploy, scale, and manage automatically compute components.

Genomes = personal profile and biometric data, e.g., a user’s voice, likeness, etc.

Governance module = wherein the settings

for a data product are decided, and the goal is to reach a consensus on the settings that must be enforced in the management module.

Human Validator App = the entry point of the validator in the Ecosystem, wherein they will share their biometric data, per industry and technology provisions on privacy established by international bodies like the European Union.

Human Validators = users who leverage their biometrics to prove that an AI is supervised.

Humans Studio = where individuals and App Builders access the AI Library and bring their ideas to life by utilizing any number of AI models and Genomes.

Inter-Blockchain Communication Protocol (IBC) = an interoperability protocol for communicating arbitrary data between arbitrary state machines.

Ingress Network Engine = an API that provides routing rules to manage external third party access to the compute engines inside the ecosystem, typically via http/https.

Load Balancing Engine = responsible for distributing network traffic across multiple compute engines.

Management module = the control panel of the person who owns the product data, where they observe how the community, i.e., the human validators, enforce the settings the owner has decided upon; the mechanics of chunking and approving the task are described here.

Model Encapsulation Oracle (HMEO) = a compute engine process where the result of a neural network is cryptographically encapsulated so that the execution of said AI model could not happen without the cryptographic signature that created the encapsulation.

Network Encapsulation Oracle (HNEO)= a network system where communication between different compute engines/oracles is fully encrypted and virtual so that for every communication between two parties, there is only one fully encrypted channel to communicate data.

Server Validators = entities securing the network and providing decentralization. They are responsible for validating transactions, user identity, data ownership, as well as ensuring all contributing parties to a valuable outcome are compensated appropriately.

NFT = Non-Fungible Tokens represent primary assets in the Ecosystem, through which we can encapsulate the digital information and maintain a persistent record of the value it produces over time.

Proof-of-Authority (PoA) = an algorithm used with blockchains that delivers comparatively fast transactions through a consensus mechanism based on identity as a stake.

Proof-of-Human (PoH) = a new blockchain consensus that relies on humans to leverage their biometric data to ensure AI's biological supervision.

Proof-of-Stake (PoS) = form of consensus mechanism that works by selecting validators in proportion to their quantity of token holdings.

Research Studio = the entry point of tech providers in the Ecosystem.

Service Discovery Engine = responsible for automatically detecting devices and services inside the ecosystem; service discovery protocols are a networking standard that accomplishes detection of networks by identifying resources.

Stakeholders = AI researchers, data providers, app builders, and human validators.

Synthetic media = also known as AI-generated media, represents the artificial production, manipulation, and modification of data and media by automated means, especially using artificial intelligence algorithms.

Validation Nodes = entities commiting new blocks in the blockchain and receiving rewards in exchange for this; they also have a role in the governance by expressing votes on new proposals, weighted according to their total stake.

zk-SNARK = "Zero-Knowledge Succinct Non-Interactive Argument of Knowledge" refers to a proof construction where one can prove possession of certain information, e.g., a secret key, without revealing that information and without any interaction between the prover and verifier.

Also, the following articles may provide references for this document:

Virtual Machine
https://en.wikipedia.org/wiki/Virtual_machine

Kubernetes
https://en.wikipedia.org/wiki/Kubernetes

Containers
https://en.wikipedia.org/wiki/Docker_(software)

Smart Contracts
https://en.wikipedia.org/wiki/Smart_contract

Blockchain Structure
https://en.wikipedia.org/wiki/Blockchain#Structure

Artificial Intelligence Network
https://en.wikipedia.org/wiki/Artificial_neural_network

Artificial Intelligence Model
https://en.wikipedia.org/wiki/Machine_learning

Distributed Computing
https://en.wikipedia.org/wiki/Distributed_computing