

Render Network Whitepaper

Originally posted August 28th, 2017 and last updated May 26th, 2023

Abstract

The Render Network is powering the rise of an open and decentralized ecosystem of connected 3D digital assets - a technological revolution that is redefining how we create, interact with, and monetize digital content. As our lives become more digital, from the surge in mobile media to developers crafting new augmented and virtual realities, our engagement with computer-generated worlds is becoming increasingly immersive. With artists and media companies embracing novel methods of content creation, the complexity faced by content creators and editors is escalating. Large-scale computation tasks and the demand for real-time virtual experiences necessitate not just additional resources, but also innovative ways to exchange digital content. The Render Network is built to provide a platform for a wide array of computation tasks - from basic rendering to machine learning training to complex calculations - which are facilitated swiftly and efficiently in a blockchain-based peer-to-peer network, free from error or delay, while ensuring secure property rights.

In order to realize the promise of next-generation immersive computing, the tools for digital creativity and innovation need to be democratized.

This is the mission of the Render Network.

Introduction

We are seeing a technological revolution reshaping how we produce, consume, and monetize media and information. A decentralized and open global rendering system is foundational for disruptive services and platforms to evolve from the next-generation world of immersive computing. In order to support the emerging media of tomorrow - ranging from cutting edge holographic worlds to AI media - it is crucial to be able to bridge the gap between the promise of these breakthroughs and the feasibility of widespread access to them. The Render Network is a [blockchain-based global rendering platform](#) where GPU work is utilized to trace digital photons and help create next generation 3D content. The Render Network processes millions of frames a year and has been used to create works ranging from blockbuster TV shows and films to large-scale stadium graphics, and finally, some of the most iconic 3D Non Fungible Token (NFT) collections. With the rise of generative Artificial Intelligence (AI), Render Network is [powering new forms of creativity](#), enabling artists to push the boundaries of immersive and interactive media. First detailed in a [2010 patent](#) for “A token based billing model for server side rendering,” the Render Network has evolved into an ecosystem for a growing range of decentralized 3D rendering and GPU computing applications.

Problem and Solution

Computational infrastructure is a major bottleneck in producing next generation media, which requires exponentially more computing power than what is available in the centralized GPU cloud. There is a competition for resources within the centralized GPU cloud - from GPU rendering and cloud streaming to AI training - leading to availability constraints and prohibitively expensive pricing for many artists.

Augmented and mixed reality content, which requires orders of magnitude more rendering power than today's HD or 4K imaging, are only exacerbating these computing shortages. As a result, artists are often burdened with costly or time consuming rendering processes when trying to create immersive imagery, limiting the potential for widespread access to next generation media production. New AI technologies are all competing for the same GPUs and unable to scale.

Meanwhile most artist's GPUs remain idle when they are not rendering their own work, or are out of use after upgrading to new GPU models, reducing the productivity potential of existing local GPU infrastructure. Further, excess GPU supply from proof-of-work cryptocurrency mining has led to an arms race where increasing compute resources are dedicated to mining fixed (or regressive) block rewards. The result: GPU's for proof-of-work blockchain mining have diminishing productivity per watt, leading many GPU models to become unsustainable - with energy consumption costs exceeding marginal revenue. With the rise of more computationally efficient Proof of Stake blockchain protocols, there is an opportunity to more productively use latent GPU compute resources.

The Render Network harnesses these unused GPU cycles by connecting creators in need of computation power for rendering their scenes or compute needs with providers with available GPU power.

Value Proposition

The Render Network makes the process of rendering and streaming intricate virtual works easier for all users. It allows complex GPU-based render jobs to be distributed and processed on a peer-to-peer network, making the transactional process of rendering and streaming 3D environments, models, and objects much simpler for end users. Node operators with idle compute power can contribute their unused GPUs to a global network and earn RNDR for processing artists rendering needs, creating a more efficient use of GPU infrastructure. Furthermore, the Render Network platform can be used for applications like crowdsourcing 3D projects to digital rights management, creating a vibrant new marketplace to fund digital ideas, assets and applications that anyone, from single developers or artists to large studios, can access and leverage.

Use Cases and Applications

The Render Network can be used by artists, engineers, designers, and developers for a wide range of GPU workloads and applications, enabling new forms of creative, scientific, and industrial production that would be cost prohibitive on local infrastructure or centralized services.

Visual Effects and Motion Graphics

With the Render Network, independent artists and studios can produce large-scale 3D visual effects which would otherwise require prohibitively costly upfront capital expenditure in on-premise render farm infrastructure or pre-reserved GPU cloud instances. Artists using the Render Network can also leverage nearly unlimited on-demand GPU resources to render scenes at higher resolutions or multiple output formats, enabling more latitude in post-production process for virtual camera movements, tracking shots, pans, and other effects.

Virtual Assets and Non Fungible Tokens (NFTs)

Using the Render Network, 3D artists can create non-fungible tokens and virtual assets with deep levels of provenance. Artists can also create and tokenize next generation virtual assets, like interactive ORBX files, enabling the monetization of portable 3D virtual assets.

Virtual Production

Directors and Cinematographers can combine live action with large-scale rendered 3D environments to transform the filmmaking process, enabling them to test unlimited takes and push the boundary of virtually created cinematic worlds.

Additional Rendering Applications

Projection Mapping

Rendering large scale 16K+ resolution 3D graphics for concert projections and venue based immersive content like stadiums or historic landmarks.

Product Design

With the Render Network's massively parallel decentralized rendering, artists can rapidly prototype at scale, testing large amounts of texture and color patterns through parallelized GPU rendering, dramatically accelerating the 3D design process.

Architecture

Through scalable decentralized rendering, architects can create immersive 3D visualizations and virtual reality renders of architectural designs, enabling clients and designers to virtually experience architectural renders and concepts with full lifelike immersion - transforming the architectural visualization process.

Simulation

Researchers can render complex physically based rendering tasks for Physics, mathematical visualizations, time-elapsed nature simulations, neuroscience simulations, and other scientific applications using the Render Network.

Visualization

Using the Render Network, engineers can visualize complex and hyperrealistic content for virtual training and complex industrial visualizations.

Emerging Applications

Augmented Reality and Gaming

The Render Network's GPU pre-rendering and real-time streaming enables artists to share real time immersive or interactive experiences for cloud streamed immersive content and augmented reality (AR). Artists stream live real time interactive 3D experiences with virtual streams using Render Network technology, or deploy hybrid pre-rendered and real time graphics to produce cutting edge cinematic gaming and AR.

Artificial Intelligence

Artists, Engineers, and Developers can create a range of AI applications using the Render Network - from prompt based AI assisted 3D content and AI accelerated holographic rendering to new forms of inference and AI training that leverage the Render Networks 3D scene graph data with full Digital Rights Management (DRM) features like attribution and traceability.

The Render Network Architecture

In its current state, the Render Network consists of two primary layers:

1. The off-chain rendering network, comprising Creators, Node Operators, the Render Network and rendering application layer vendors. Node Operators are GPU nodes that provide power to the network.
2. Blockchain layers that handle payments, via RNDR and escrow contracts. By leveraging the blockchains' public ledgers, all Creator-Node Operator interactions are publicly verifiable, giving Creators, Node Operators and the Foundation Team the ability to ensure all transactions are processed correctly, and if not, are easy to track and correct.

Node Operators and Creators interact with each other in a distributed network model that is arbitrated by a core network infrastructure of servers.

OctaneBench is a proprietary benchmarking tool (and the most popular benchmark for GPU rendering used today) that measures GPU rendering speeds, measured in OctaneBench Points per Hour (OBh), which in turn determines the pricing of rendering tasks. OctaneBench is used to standardize and benchmark a GPU's performance, ensuring that Creators are subject to consistent performance-based pricing when requesting work from Node Operators. The cost for any task is measured in OctaneBench and is determined by a [Multi-Tier Pricing \(MTP\)](#) algorithm.

Reputation Scoring

As Node Operators successfully complete jobs on the network, they build trust needed to receive a higher volume of jobs. Node Operator reputation is built by timely and accurately completing jobs. Creators build reputation scores by successfully using the network. As creators build a history of successful jobs - with minimized user error - they are able to access larger amounts of concurrent GPU nodes. The reputation score, therefore, helps the network efficiently assign work and reduce unintentional congestion from failed renders or malicious [Sybil Attacks](#). The reputation scoring system is periodically updated to increase network efficiency as the complexity of work and the network's service offerings evolve over time.

Allocation of Resources

Resource allocation on the Render Network can be broken down into 2 categories: Job Allocation and Node Allocation, which function as follows:

Job Allocation

Job Allocation is the process by which a Creator's scenes are allocated to Nodes on the network. Currently the Render Network's job allocation system prioritizes a Creator's reputation score alongside their scene size and concurrent nodes available at that time.

The creator's reputation score factors into their job allocation priority in 2 ways:

1. Creators with higher reputation scores can run more concurrent jobs on the network than users with lower scores.
2. A priority modifier is added onto each Job based around the Creator's current reputation score, making Creator's jobs who have high reputation scores 'more appealing' to nodes than equivalent jobs by lower-reputation score Creators.

Node Allocation

Node Allocation is the process by which Nodes are assigned rendering work on individual jobs. When jobs are sent to the Render Network, they are assigned to node operators based on factors such as the selected tier, hardware requirements, time on the network, user reputation, and node reputation score, and OctaneBench score.

Privacy and Security

The Render Network is built to handle tasks for a wide range of users, from individual artists to globally recognized studios. When rendering or working with any creative or confidential material, privacy and security is of extreme importance. In order to protect the privacy of Creators and assets in the network, the system uses a combination of end-to-end encryption, hashing, virtualization and secure centralized storage:

- Every scene is broken into many individual assets, which are hashed and encrypted while uploaded to the network.
- Rendered outputs are encrypted before being sent through the network.
- Any asset stored in memory or on disc is always encrypted.
- All individual frames are watermarked prior to download, ensuring payment is given before scenes are downloaded.
- Individual asset storage is short term.

The Render Network API and SDK

The Render Network provides an API, enabling users to build commercial third party applications on top of the Render Network. The Render Network API supports customizable asset rendering pipelines, for example swapping specific files like textures on the Render

Network, rather than re-uploading complete ORBX files. This enables users to build new micro-services and applications ranging from automated batch rendering asset creation pipelines to next generation real time streaming and generative art production.

Using the Render Network API, administrators can create AUTH permissions for multiple users in a shared studio, with credential management tools for uploading scenes and submitting jobs. As a result, a studio can customize their account administration. An SDK will be made available which will enable developers to create and integrate software applications to the Render Network.

The ORBX File Format and Standard

The Render Network uses the [open ORBX file format and streaming framework](#) in order to support fully distributed rendering. ORBX is a container format that captures scene data (assets) and an XML (extensible markup language) render graph which describes the semantics of a scene. Just like a web page, it can be cached to archive (.orbx file), or streamed from a URL or URI over raw UDP/tcp or web wss or https. The ORBX format supports over 20 of the industry leading DCC (Digital Content Creation) tools, and contains industry standard sub formats like Alembic, OpenVDB, EXR, Open Shader Language (OSL), and glTF. ORBX exporting is increasingly becoming natively integrated into the OctaneRender interface [enabling simplified one-click exporting processes](#) to the Render Network from a host of third party plugins alongside improvements in exporting speeds and compatibility with expanded third party plugin effects and libraries.

The ORBX interchange format is critical for abstract rendering work from host 3D applications. A one-click ORBX export from a 3D content creation tool fully decouples all assets and code needed to perform a remote GPU render job on multiple mining nodes. By fully abstracting scene data from third party software tools with the ORBX scene format, the Render Network is able to parallelize work across a blockchain peer-to-peer network at near unlimited scale without dependencies on local software. The ORBX interchange works regardless of host application, providing efficient, open, software-agnostic distributed rendering.

The Render Network also leverages the ORBX scene graph to enable deep chain of authorship and validation. Every time a user uploads a scene and a node operator processes a job on the Render network, all assets and settings in the ORBX render graph are hashed. With assets in a scene attributed in ORBX XML render graph, the network has a semantic history of every object and setting within a scene. The hashing of the ORBX render graph and scene data provides immutable and granular history of all assets and work processed on the Render Network. Thus, as assets move through the network, ORBX provides full traceability needed for attribution and authorship.

ORBX has recently become the basis for the [ITMF \(Immersive Technology Media Format\)](#), an open-source file standard for emerging holographic and 6dof immersive 3D content, and is being leveraged by a number of web3 and metaverse related working groups as a standard for next generation blockchain media distribution. The format enables interoperability across 3D tools as well as the distribution of fully volumetric six-degrees-of-freedom (6-DOF) scene data for fully immersive holographic experiences and virtual assets. ORBX.js streaming technology

delivers high performance 3D games and desktop applications to the open Web – using only HTML5 and JavaScript. Because ORBX.js is browser based, it bypasses apps or downloads, enabling frictionless publishing of immersive media experiences to multiple endpoints like VR, AR, and mobile. The ORBX framework is open and accessible in OctaneRender Render Modules, allowing third party developers to build additional ORBX-based applications or services through an API.

Both ORBX and Octane are extensible, supporting most mainstream 3D content creation tools. Capabilities like Delta sync enable users to leverage the Render Network to precompute scene changes without re-uploading a new ORBX file.

Multi-Render Support

As part of the decentralization of the network’s architecture to accommodate more artists’ 3D workflows, support for new third party render engines will be added to the network. Integration of Maxon’s Redshift is in development with Arnold Render and other emerging open source engines like Blender’s Cycles to follow.

Note, thirty party render engines are entitled to charge license fees, and if they choose to do so, these will be added to the fee charged for a job. OTOY currently waives such fees for OctaneRender.

The Render Network and AI

The growth of artificial intelligence (AI) has generated an unprecedented demand for computational resources. AI applications necessitate vast computational power, outstripping the capacity of traditional CPU-based systems. Through the Render Network SDK, developers will be able to leverage the network’s decentralized GPUs for AI compute tasks ranging from NeRF (Neural Reflectance Field) and LightField rendering processes to generative AI tasks.

Increasingly, 3D artists are introducing AI generated content into their creative workflows, combining hand created digital artwork with generative AI processing. With the integration of AI toolsets like Stable Diffusion on the Render Network, the network supports the increasing convergence between traditional and next generation creative workflows that leverage AI. For example, artists can use artificial intelligence tools to create assets like generative AI textures that are used to render ultra-high resolution immersive 3D worlds on the network. Large-scale art collections using generative AI to vary outputs can also be distributed across the network’s nodes, enabling creators to frictionlessly create AI art collections at near unlimited scale. The rise of artificial intelligence requires new forms of digital traceability and asset verification. The Render Network’s deep levels of on-chain provenance built into each work’s Render Graph enable licensing 3D models for AI training, or royalty based usage.

The Render Network uses AI technology to accelerate and optimize rendering processes. AI denoising in OctaneRender, the engine used in the Render Network client, has been specifically trained to denoise volumes with further training optimizations possible through distributed computing. Scene AI models surface visibility to get maximum speed while denoising and

rendering Out-of-Core Geometry and Emissive Objects. These models are trained on perceptual models of Material, Spectra Irradiance, and Scene Data and enable accelerated rendering for more complex scenes when scene data exceeds VRAM capacity. These models are periodically updated and can be further trained using decentralized GPU nodes.

The Render Network and Virtual Assets

All work on the Render Network is hashed, including the ORBX file an artist uploads to the Render Network, each individual frame from a render job, and an animation produced from the completed frames. The hashing data associated with a render job, ORBX file, or individual frames can be minted on chain and included in NFT metadata.

In addition to completed render jobs and frames, all scenes uploaded to the Render Network are hashed with representation of a scene's XML data as well as each of the individual assets contained within a scene. The unique hashed IDs for each scene and its assets can be minted in on-chain metadata, providing additional levels of provenance for blockchain assets.

The hashing of each scene graph and its composition of assets enables the network to track and represent changes to the state of each scene graph over time through a process called delta syncing. As blockchain virtual assets and NFTs transition from static images and video to dynamic or interactive works of art that change over time, delta syncing of hashed on chain scene graphs enables the tokenization of these real-time experiences.

Decentralized GPU Streaming Experiences

Art produced on the Render Network can be streamed using the X.IO GPU application streaming framework, enabling the distribution, consumption, and monetization of ultra high resolution media directly in-browser with no downloads or hardware dependencies. The technology for GPU streaming is available at x.io, and GPU streaming applications running on decentralized Render Network nodes are in development to meet user needs for scalable 3D applications, services, and experiences that can be monetized on a pay-as-you-go basis using RNDR.

Job Procurement and Payment

Multi-Tier Pricing (MTP)

To make the network as flexible as possible, Multi-Tier Pricing (MTP) is available to the network. Creators choose from multiple pricing tiers according to their preferences for speed, cost, security, and Operator reputation.

Tier 1 Service is the most expensive and thus doesn't benefit from any pricing discount multiplier (i.e., the Tier 1 multiplier is always 1). Tier 1 service utilizes a separate pool of verified

nodes and jobs are considered to be the highest priority with access to GPUs that have higher VRAM for high complexity scenes.

Tier 2 Service represents the highest tier of decentralized rendering, with access to the highest performance and most parallel GPU nodes. When compared to Tier 1, Tier 2 benefits from a pricing discount multiplier that ranges from 2 to 4, depending on network conditions.

Tier 3 Service still tries to ensure full job completion, jobs are slower to process and operating nodes may have less advanced GPU hardware than those in Tier 2. Tier 3 service has the highest discount multiplier, typically ranging from 8 to 16, depending on network conditions.

Expansions of the tier and pricing structure can be made dynamically based on changing supply and demand conditions, including introducing specialized or high performance tiers, AI computing tiers with AI hardware, tiers with renewable energy sources, and other changes as the preferences of network users evolve over time.

Burn and Mint Equilibrium (BME)

In [RNP-001](#) the community voted to implement the Burn and Mint Equilibrium (BME) model for managing how the network operates.

The BME enables the network to consistently price services within a multi-sided decentralized economy, allowing Creators to predictably manage their rendering costs and Node Operators to provide rendering services. In the BME, rendering and other forms of work on the network are priced in fiat. Pricing in fiat does not preclude paying in an equivalent amount of RNDR.

Once the work is completed and the proof of render is accepted on the blockchain, Creators use 95% of their fiat fee to programmatically purchase RNDR from distributed liquidity pools, and to subsequently burn it. A public log of all these transactions is used to calculate proportionate emissions-based rewards for contributors at the end of each epoch. The remaining 5% is paid to the Foundation as a transaction fee in order to fund Render Network operations.


Separately, emissions are used in the BME to incentivize activity that enables the development of a multi-sided decentralized network bringing together supply (Node Operators), demand (Creators), and infrastructure providers (Liquidity Providers). In addition to matching supply and demand, network functioning requires forming and incentivizing distributed liquidity pools to allow for sufficient liquidity to facilitate automated RNDR purchases.

Emissions follow a preset and declining schedule that can be viewed [here](#), and please note under [RNP 003](#) 4.56M RNDR of the first year emissions will be subsequently removed from first year emissions and allocated to the Render Network Foundation.

Emissions will be allocated on an epoch by epoch basis. Epochs occur based on use of the network, generally a 24 hour period and can be adjusted by the Foundation with input from

Governance as needed in order to respond to changes in network usage.

RNDR Burn and Mint Equilibrium



Emissions Incentives

Emission distribution per epoch will be adjusted as per the growth requirements of the network i.e. valued stakeholders at different stages could receive greater shares of emissions. If the system is running near equilibrium state, the fulfillers will always be paid the appropriate amount.

Here is how the Render Network burn-and-mint model could work for Creators, Node Operators and Liquidity Providers:

1. **Creators:** On an epoch basis, Creators could receive a percentage of their RNDR spent that epoch back in the form of RNDR. This could ideally incentivize further use of the network and reward power users by allowing for more creations to be rendered on the Render Network. Percentage returns could be as high as 100% of RNDR spent initially, and could gradually taper as time goes on.
2. **Node Operators:** Node operators will be rewarded for performing work or providing value. These rewards could generally be segmented into two buckets:
 - i. **Availability Rewards:** Incentives for node operator liveness.
 - ii. **Job-to-be-done completion rewards:** Incentives for completing jobs submitted to the network.
3. **Liquidity Providers:** Liquidity providers will be rewarded per epoch for contributing staked RNDR to the liquidity pools on partnered exchanges, allowing for RNDR to be available for the new Burn and Mint Equilibrium system. Liquidity providers could be rewarded with an additional percentage of RNDR staked in that epoch, as determined by the Foundation.

Net Emissions Cap

A net emissions cap will be designed to ensure that in the long term, once the supply cap for the base RNDR has been reached by the network, fulfillers are still able to receive rewards for

performing work on the network. This will be managed via the RNP process. When implemented, it will work as follows:

1. On an epoch by epoch basis, net emissions will “recycle” some set of burned RNDR available for use in rewarding fulfillers when there are no longer enough provided through the scheduled Emissions.
2. Net emissions allow rewards to exist in later years, however, they will need to be capped at some percentage of current issuance to maintain the core BME.

Implementation of the BME

The network currently consists of two key stakeholders: Creators (or parties with jobs-to-be-done) and Node Operators (or parties performing jobs-to-be-done).

Example: The job-to-be-done is “provide rendering power for creators” over the Render Network. The requester is “creators looking to get their projects rendered on the Render Network”. The fulfiller is the actual nodes processing these jobs.

The proprietary on-chain record of transaction activity is dominated in RNDR.

Each job-to-be-done would be priced in fiat, and the requestor has the option to pay in either fiat or the equivalent amount of RNDR. Once the requestor approves the work performed, the protocol purchases and burns the required amount of RNDR (determined by market prices for RNDR at the time of the approval), and records the activity in a public log.

Note that the amount of RNDR burned to access the underlying service would be denominated in fiat based on the transaction price, but if the creator/customer elected to pay in RNDR, the denomination of the RNDR burnt would be the RNDR at the time of entering into the transaction, not at the time of confirming completion of the work.

The intended effect of this is that the requestor would demonstrate on-chain that the fulfiller has completed the work for the money that was burned.

Example: A Render Network creator burns \$5 worth of RNDR to produce \$5 worth of RNDR Job Credit, which are non-fungible and non-transferrable and used purely to track the relative percentage of job-to-be-done completion rewards allocated.

Node operators would be compensated for completing jobs via RNDR issuance incentives:

Independent of the RNDR burning and RNDR Job Credit issuance process, the protocol would emit a number of RNDR per epoch, set by the Foundation and Governance, and distribute them to fulfillers according to predefined rules.

Node operators would therefore be compensated for performing work or providing value on the network in the base RNDR asset. These rewards could generally be segmented into two buckets;

- a. Availability rewards and incentives: incentives for showing that the fulfiller is available to complete work on the network
- b. Job-to-be-done completion rewards: incentives for actually completing jobs-to-be-done submitted to the network

Example:

Job-to-be-done completion rewards: If 2% of RNDR burned during an epoch were in the name of Node Operator A as represented in the Render Job Credits log, then Node Operator A would receive 2% of Emissions allocated for job-to-be-done completion rewards.

Availability rewards: If the same node were to complete 5% of uptime challenges defined to assess the reliability of the node network, it would receive an additional 5% of Emissions allocated for availability rewards.

Rendering Engine Fees

The Render Network protocol will charge a 5% network fee on all transactions, which is imposed to cover any infrastructure costs for operating the network, including continued employment and eventual expansion of the Render Network Foundation team.

RNDR Information


RNDR was originally created on the Ethereum blockchain.

In 2018, RNDR migrated to a new contract, decreasing the supply at a 4:1 ratio and introducing additional smart-contract functionality, arriving at a max supply of [536,870,912 RNDR](#), (or 2^{29}). Migration of RNDR to the new contract is open indefinitely. Instructions for migration can be viewed [here](#). The RNDR contract is [0x6de037ef9ad2725eb40118bb1702ebb27e4aeb24](#).

In January 2023, the [Render Network community voted to approve RNP-001](#), which introduced an emissions mechanism to administer the Burn-Mint Equilibrium (BME). [RNP-001 proposed a 20% inflation pool](#) that is released as network incentives to Node Operators, Creators, Liquidity Providers, and Consumers to advance the functioning of a multi-sided network. With the passage of RNP-001, the supply of RNDR will increase by 107,374,182 over the course of two periods (years 1-5, or a launch period) and (years 5 to perpetuity) reaching a maximum theoretical supply of 644,245,094 (prior to any RNDR burned as outlined under the BME model), with initial emissions modeling outlined in RNP-001 as follows.

Total Supply	644,245,094
Pre-Inflation Supply	536,870,912
Inflation Amount	107,374,182

RNDR Emissions



The specified schedule only defines a cap on the total circulating RNDR, and does not account for the substantial amount of burn that will occur as a result of network activities. The increase in supply is released in small batches (never more than 10% in a given year) and dilutes existing holders in favor of new entrants into the network.

The emissions schedule is broken down into two periods: Launch (Years 1-5) and Growth (Year 5 - Indefinite). The purpose of this is to design a front-loaded rewards schedule in initial stages of the network. In the growth stage, RNP-001 defines gradually decreasing outflows as per the damping coefficient (proposed at 0.945).

The emissions schedule outlined in RNP-001 is subject to forthcoming governance procedures, and relies heavily on parameter tuning in order to reward actors in later stages of the network. The protocols specified in the Render Network Proposal (RNP) governance process, and passed in RNP-000, can inform changes to the schedule.

With the passage of [RNP-002: Layer 1 Network Expansion](#), the implementation of the BME will take place on Solana, requiring migration and a new contract on the Solana blockchain. Information about this process will be detailed in forthcoming governance proposals and updates that can be viewed on renderfoundation.com.

Architecture Improvements

Improved Reputation Algorithms and Reputation Score Mapping

The rendering network relays critical data to the public blockchain layer. Specifically, any task between Creators and Node Operators that is requested is time-stamped and recorded on-chain along with asset hashes. As Node Operators complete tasks, Creators are able to review and check rendered frames and thus, the performance of the Operator, with the result being factored into an Operator's dynamic reputation. Likewise, Creators are also subject to dynamic reputation scoring, based on whether they approve work that has been properly rendered. Actor reputation scores form the basis of Render Network's Proof of Render reputation system.

Improvements in Order-Matching

In order to effectively match GPU requests to providers, the Render Network incorporates node reputation history and node power in its automated assignment process. Using Multi-Tier Pricing, creators select from a menu of preferences for cost, speed, and security, enabling the network to optimally sort jobs based on the nature of demand at any given moment. The assignment process incentivizes GPU providers to maintain a high node success rate and allows creators to optimize their preferences. Job assignment and Multi-tier pricing are based on a tier system.

Improvements in Pricing

Decisions concerning pricing greatly affect whether the network can fairly compensate both the Render Network team and GPU Providers for their respective contributions (network infrastructure, proprietary software, rendering performed, etc.) to Creator rendering services. Pricing decisions can also affect network growth and competitiveness.

Time-Based Billing

In order to deliver pure AI training and inference jobs, and potentially decentralized GPU streaming experiences, the billing model may need to be expanded to allow time-based rentals necessary to secure the required dedicated hardware.

Dispute Resolution Mechanisms

If a creator submits a job and a node delivers subpar work, the creator can initiate a dispute resolution process. Eventually, a review committee will be charged with reviewing the submission and slashing the node operator in case they have been at fault.

Ownership and Governance

RNDR mediated incentives and governance through the Render Network Protocol (RNPs) help a decentralized set of participants reach an equilibrium of a functional utility.

The Render Network Foundation

The Render Network Foundation is a not-for-profit dedicated to maintaining the core Render Network protocol and growing its community. The Foundation facilitates the governance process through the Render Network proposal (RNP) system. It also is responsible for helping to set the strategic priorities of the network and issuing grants to support them.

Render Network Protocol

The Render Network will be ever changing based on community needs and input through the process of Render Network Protocol (RNP). [The RNP-000](#) was implemented to define how these changes take place. RNPs are the process by which the community can affect change on the Render Network. Outlined below are all of the steps through which a community member can go from idea to implementation onto the Render Network.

1. Initial Proposal
2. Draft Submission
3. Initial Proposal Vote
4. Render Network Team Review
5. RNP Vote
6. Implementation

All RNPs status will be listed on the [Render Network Github](#).

Twitter: <https://twitter.com/rendertoken>

Discord: <https://discord.gg/rendertoken>

Telegram: <https://t.me/s/rendertoken>

Medium: <https://medium.com/@rendertoken>

Github: <https://github.com/mdr-network?tab=repositories>

Snapshot: <https://snapshot.org/#/rndrnetwork.eth>

Foundation: renderfoundation.com

Website: <https://render.x.io>