Manifest AI

White Paper

Organized by The Lifted Initiative

Bravick, Eric
Primary Author/Inventor
ebravick@liftedinit.org

Morency, Felix
Primary Maintainer
felix@liftedinit.org

Onyeka, Obi
Legacy Author

Sloan, Tyler
Legacy Author

October 7, 2025

1 Introduction

The Lifted Initiative is on a mission to break down the barriers to Web3 mass adoption through building Manifest, a Layer 0 network that is interoperable across generations of the web. Layer 0 networks such as Polkadot [8], Avalanche [7], and Cosmos [5] scale with interoperable sub-networks. The Manifest Network is able to take this a step further by bridging Web 2.0 and Web3 tech via a full stack approach that enables blockchain and non-blockchain sub-networks to communicate. Manifest provides an easy way for today's companies and developers to gain the benefits of decentralized technology without having to hire expensive additional resources or learn a new programming language.

This paper details the vision and function of the Manifest Network Token economy. This white paper is not intended to be a complete technical specification, and its scope is specifically restricted to the design and function of the token economy as an economic construct. This paper does not discuss the technical architecture of the network, framework, or protocol that collectively produce the Manifest Network, or their related technical components.

The Manifest Network provides computational services in a cryptographically verifiable decentralized network. Its unique design necessitates a native token, MFX, which, among other functions, incentivizes validators on the network, enables decentralized validation actions taken on the network, and enables decentralized governance of the network by the holders of the Token.

Additionally, the network makes use of PWR, which is a digital credit that is redeemable for services running on the Manifest Network (e.g. CPU, Memory, Storage, etc). This credit is a discrete feature in the Manifest Network. While PWR is a cryptographically-secured mathematical representation of the ability of the holder to use the network, PWR is not intended to function as a store of value or medium of exchange. PWR, therefore, is not intended as a "crypto" asset with monetary properties, but as a cryptographic credit representing resource reservations.

The architecture of PWR enables the Manifest Network to price decentralized services based on the production cost of those services, rather than basing prices on a market-driven token price (e.g. BTC, ETH, etc.). Furthermore, PWR allows for a seamless experience for the developer and end-user, reducing barriers to network adoption by reintroducing classic Web 2.0 cost dynamics into a Web3 style network.

1.1 Background

Many token economies correlate the cost of using network services with demand for the network token, which can result in over-pricing the services provided by the network.

Single asset token economies services priced in a self-regulating manner, such as Bitcoin [6] and Ethereum PoW [3], tend to lack overall utility. These assets rely on the supply and demand curve of the token itself to determine a user's cost of doing an operation on that network. Consequently, the cost of operations, such as transactions, are highly volatile and subject to enormous increases in cost [4]. This correlation of the cost of network operations with token price, rather than production price, has prevented these networks from being used for real utility purposes.

In contrast, the supply and demand curve of the US Dollar is generally uncorrelated with the cost of cloud computing. The cost of cloud computing is correlated with the cost of producing the cloud compute resource, including the cost of pricing servers, bandwidth pricing, etc.

Infrastructure networks should continually incentivize new stakeholders to join a network to provide services and support to the ecosystem as a whole [2]. Successful production-oriented networks must correlate the cost of using network services with the production cost, not token price, especially when working with business owners relying on predictable and stable costs. Token economies should scale with the size of the decentralized system and be balanced with consumer demand to reach equilibrium via an adaptable approach [2].

The Manifest Network's token economy is a burn-and-mint equilibrium (BME) [1] design that uses a dual asset model to correlate the cost of network operations with the production cost, while maintaining the many advantages of a network store-of-value utility token that may appreciate in price. With a BME design, a central reserve "mints" tokens to reward suppliers and "burns" tokens when consumers want to utilize network services with the overall goal of maintaining a stable supply-demand equilibrium [2].

2 Tokenomics Model

The Manifest Network tokenomics model is based on a dual-asset model: MFX and PWR. MFX is a commodity token that stores value on the network, facilitate governance and ensure proper network security. MFX is also used to rewards participants that provides work on the network, such as ... PWR is a utility token that is redeemable for network services. The dual-asset model goal is to correlate network usage costs with production costs rather than with the token price.

Participants in the manifest network can carry any amount of manifest token (MFX), power token (PWR), and/or any amount of a finite number of other assets. We assume that each participant has a utility function

$$U(a_{MFX}, a_{PWR}, a_1, a_2, ...) = \ln a_{MFX} + \ln a_{PWR} + \sum_i \ln a_i$$

that is a function of the collection of assets a_i that each participant contains at some point in time. The logarithmic utility model is used to best model diminishing marginal utility as the quantity of assets accumulated increase, which is typically the case for participants that have aversion to risks associated with accumulating those assets. The utility function takes values in some totally ordered set. Each participant has the ability to perform an action that, in turn, changes the distribution of assets that they have, and, therefore, changes the value of that participant's utility function. We are interested characterizing the set of participants and actions for which the utility functions of all participants are maximized. Sets of participants and actions that do this align incentives among all of the network participants, and so we call them incentive equilibria. In particular, we are interested in equilibria that include actions taken that the Manifest Network specifically provides, thereby justifying the construction of the network itself.

Given a certain set of boundary conditions, it could be the case that the participant utility function, as it has been defined above, has a closed-form solution. However, for the purpose of determining what happens in equilibrium, it is not necessary here, due to the fact that we previously restrict participants to a finite number of actions. Therefore, the primary use of the utility function for determining equilibrium in this setting is for choosing which of those actions a participant will perform next. Since, at any given point in time, the choice of action comes from the same set, we can describe an equilibrium set of actions for each participant without having to derive a closed-form utility maximization.

2.1 Participant Actions

In this model, participants can take the following actions:

- Send an amount of MFX to another participant
  - This action is essentially a gift, and so will not be taken into account in utility maximization.
- Burn the spot rate of MFX to acquire PWR
  - The marginal utility for this action can be calculated as

    $$lnSR_{\frac{PWR}{MFX}}$$

    where $SR_{\frac{PWR}{MFX}}$ is the spot conversion rate for MFX to PWR.
- Stake an amount of MFX to act as a node resource provider for other participants (node provider)
  - In staking, a participant relinquishes an amount of MFX in the present, only to receive the same amount, along with an additional reward, in the future. This roughly calculates to a marginal utility of

    $$\ln d_t * (b_{MFX} - a_{MFX})$$

    where d_t is a discount factor for MFX values at a future time t, a_{MFX} is an amount of MFX equal to the opportunity cost of what would have been earned, had the amount of MFX staked been instead converted to the risk-free asset a_{rf} and held in savings, and b_{MFX} is the future MFX reward.

- Send PWR to a node provider in exchange for a resource
  - Here, we model the exchange of PWR for resources (compute, storage, bandwidth) indirectly, by considering the resources consume to yield an asset a_i taken into account in the participant's utility function. This calculates to

    $$\ln SR_{\frac{a_i}{PWR}}$$

    , where $SR_{\frac{a_i}{PWR}}$ is the spot conversion rate for PWR to resource a_i.
- Stake an amount of MFX to vote on a change a parameter of the network (governor node)
  - Token conversion rate
  - Transaction fee amount
  - Utility costs in PWR
    * Compute
    * Storage
    * Bandwidth

  Staking here works the same as in staking as a node provider.
- Stake an amount of MFX to vote on the next block of the blockchain (validator node)
  - As with the prior staking action, the marginal utility of this staking action involves calculating the difference between a future amount gained as a reward and a future fractional amount of MFX already owned being lost, by virtue of staking a fixed amount in the present and receiving that same nominal amount in the future.

2.2 Model Constraints

- No arbitrage in staking actions
  - If node provider staking is more lucrative, blocks do not advance and conversion rates become eternally constant.
  - If validator staking is more lucrative, participants cannot request PWR and, therefore, cannot access resources that the network provides.
  - If governor staking is more lucrative, blocks do not advance and participants cannot access network resources.
- No arbitrage preferring staking to exchange, or vice versa
  - A preference for staking over exchange kills commerce and defeats the utility of a multi-token network. A preference for exchange of staking is not viable due to the necessity of staking for advancing blocks and governing exchange rates.

- No arbitrage with respect to time
  - Participants can expect to earn a nonzero risk-free rate of at least one asset a_{rf} that is meant to model the equivalent of a modest savings rate. This is to ensure that there is no advantage to saving versus investing assets.

2.3 Equilibrium Conditions

In this section, we explore implications of the arbitrage conditions in the previous section, as well as make note any interesting non-trivial equilibria (equilibria where the network is leveraged).

Setting the prior model constraint of no arbitrage between staking actions leads us to model staking marginal utility as a single equation. We use this equation to explore the implication of the second model constraint, that of there being no preference between staking and exchange.

$$SR_{\frac{PWR}{MFX}} = d_t * (b_{MFX} - a_{MFX})$$

This equation is the result of setting the marginal change in utility for exchanging MFX for PWR with the marginal change in staking MFX. It suggests that, for proper network function, governor nodes must NOT be allowed to choose *both* the spot conversion rate between MFX and another asset and the reward for staking MFX, since knowing one and adhering to the no-arbitrage condition implies a calculation to retrieve the other.

Now knowing the conditions under which the actions available to each participant in the Manifest network do not lead to an opportunity for arbitrage, the next question is whether or not it can be said that there is a preference for participating in the network versus not participating. This is the same as a participant asking whether there is a benefit, with respect to the participant's utility function, of participating in the network.

A participant of the network can only exchange one asset or another or stake for MFX. Assume that, without aid of the network, each participant experiences a marginal utility (gain), dependent on some external asset a_{ext}, of $dU(a_{ext})$. Since one can only stake for MFX on the network, it can only be the case for that participant that using the network is beneficial if the participant experiences a greater marginal utility utilizing the network for exchange than the marginal utility associated with the external asset. That is, either $\ln SR_{\frac{a_{ext}}{MFX}} > dU(a_{ext})$ or $\ln SR_{\frac{a_{ext}}{PWR}} > dU(a_{ext})$. This translates to saying that burning PWR to acquire an asset that has utility to a participant is more efficient than another method of acquiring that same asset.

3 (Provisional 10/6) Conversion Rate Determination Between MFX and PWR

The conversion rate between MFX and PWR is defined to track the market exchange rate implied by verifiable on-chain activity and validator-reported quotes. Formally, the real-time conversion rate $SR_{\frac{PWR}{MFX}}$ maps units of MFX irreversibly burned into units of PWR minted, aligning the burn–mint process with observable market conditions and the cost of providing network services.

3.1 Market-Based Conversion Framework

Currently, without staking, PWR is converted strictly on a market-based conversion framework. At time t, the spot conversion rate is specified as:

$$SR_{\frac{PWR}{MFX}}(t) \;=\; \frac{P^{\text{ref}}_{MFX}(t)}{C_{PWR}(t)},$$

where:

- $P^{\text{ref}}_{MFX}(t)$ is the reference market price of MFX, derived from decentralized exchange (DEX) data (e.g., time-weighted or volume-weighted prices across eligible pools) and corroborated by validator node price submissions.
- $C_{PWR}(t)$ is the service-credit index representing the cost-normalized work encapsulated by one unit of PWR (compute, memory, storage), expressed in the same numeraire as $P^{\text{ref}}_{MFX}(t)$.

This structure ensures that PWR issuance is economically neutral with respect to both external token prices and internal production costs, while avoiding reliance on discretionary parameters.

3.2 Reference Price Construction and Smoothing

To enhance robustness, $P^{\text{ref}}_{MFX}(t)$ will be computed from a basket of admissible liquidity pools and trading venues using a time-weighted average price (TWAP) or volume-weighted average price (VWAP) over a sliding window W:

$$P^{\text{ref}}_{MFX}(t) \;=\; \text{Median}\Big(\text{VWAP}_i(t; W)\Big)_{i\in\mathcal{V}},$$

with $\mathcal{V}$ the set of eligible venues/pools that meet minimum-liquidity and data-quality thresholds. Validator-reported quotes serve as an independent cross-check and are aggregated via robust statistics (e.g., median-of-means) to mitigate outliers or transient market microstructure noise.

A protocol-level exponential moving average can be applied to reduce short-horizon variance without distorting price discovery:

$$\widehat{P}^{\text{ref}}_{MFX}(t) = \lambda\,\widehat{P}^{\text{ref}}_{MFX}(t-\Delta t) + (1-\lambda)\,P^{\text{ref}}_{MFX}(t), \quad \lambda \in [0,1),$$

and the smoothed reference $\widehat{P}^{\text{ref}}_{MFX}(t)$ may be substituted for $P^{\text{ref}}_{MFX}(t)$ in the main formula when volatility is elevated.

3.3 Operational Safeguards (Non-Discretionary)

To preserve integrity and resist manipulation:

1. **Venue Eligibility:** Only venues/pools exceeding defined liquidity, age, and uptime thresholds contribute to $P^{\text{ref}}_{MFX}(t)$.

2. **Outlier Rejection:** Prices deviating beyond a fixed multiple of recent dispersion (e.g., median absolute deviation) are excluded from the aggregation set.

3. **Volatility Controls:** If the relative change $|\Delta P^{\text{ref}}_{MFX}/P^{\text{ref}}_{MFX}|$ exceeds a preset bound within W, the system enlarges the lookback window or temporarily favors the smoothed estimator $\widehat{P}^{\text{ref}}_{MFX}(t)$ for conversions.

4. **Atomicity:** Conversion is executed as a single, atomic operation that burns MFX and mints PWR in the same state transition, preventing partial execution.

3.4 Economic Implications and Current Status

By tying $SR_{\frac{PWR}{MFX}}(t)$ to an externally verifiable reference price and an internally measured service-credit index, the conversion mechanism:

1. Aligns PWR issuance with compute and storage provisioning costs rather than speculative fluctuations alone;

2. Preserves predictable pricing for end users consuming network resources; and

3. Ensures node operators are compensated in proportion to real resource utilization.

Until the on-chain conversion module is enabled, the same formula is applied off-chain to quote equivalent rates for operational workflows, with sources and calculations recorded for auditability. Once activated on-chain, all inputs (DEX prices, validator quotes, and index values) and the resulting $SR_{\frac{PWR}{MFX}}(t)$ will be emitted as events to facilitate independent verification and indexing.

3.5 Simulation

We have been working with a research team at the University of Texas at Austin to tweak their tokenomics dynamics product for the purpose of using it to produce simulations of Manifest network participants on the network and the benefit these participants receive from being on the network. While this simulation effort is still considered a work-in-progress, we have been able to produce a simulation that illustrates a participant gaining utility from the network using the assumptions of our underlying mathematical model.

```json
{
    "profile_type"    : "MICRO",
    "node_types" : [
        {
            "name": "Bias Token 1",
            "number": 100,
            "utility_variance": 1,
            "utility_matrix": [6, 3, 2],
            "utility_threshold": [0.7, 0.7, 0.7]
        },
        {
            "number": 100,
            "name": "Bias Token 3",
            "utility_variance": 1,
            "utility_matrix": [1.2, 2, 3.1],
            "utility_threshold": [0.7, 0.7, 0.7]
        }
    ],
    "token_types" : {
        "num_tokens": 3
    },
    "initial_state" : {
        "token_supply": [100.0, 10, 5]
    },
    "simulation_length": 100
}
```

The tokenomics dynamics product is driven by a configuration file that allows one to list the configuration of each participant. Like in the mathematical model, each participant is given a utility vector, the entries of which represents a log-linear factor with which that participant's utility increases for each unit of that token type acquired. For simplicity, each token type in the simulation is assigned an index, but given no other identifiable traits. An initial state configuration allows for setting of the initial token supply of each token type. In the Manifest network, there is only one utility token, which is PWR, but we produce simulations using three token types to represent one each of three different *types* of resources that might be consumed by burning PWR (that being compute, storage, and bandwidth).

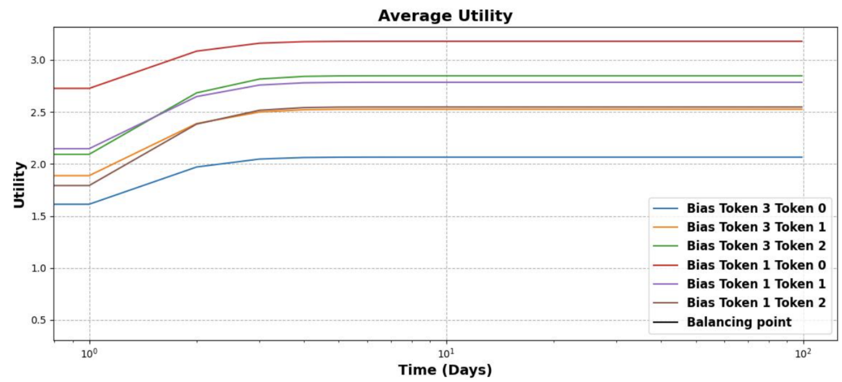

In the simulation illustrated, we assume three token types, provide two participants, and give each participant and positive utility bias toward a different token type. The simulation shows that the token type bias affects the rate at which tokens of that type are accumulated, but, ultimately, net average gains in utility are shown across all token types for each participant, up until a plateau is reached. The simulation appears to summarize a short burst of utility gain by utilizing the network, before reaching a steady state. We are, so far, unable to alter the timeframe in which a participant moves to and achieves this equilibrium, and that is something worth further investigation with the simulation tool.

4 Dual Asset Model

The goal in designing this dual asset model is to make the Manifest Network more desirable for individuals and enterprises to conduct business by preventing volatility in production costs. Manifest Networks' token economy design has two distinct features: the MFX cycle, described on page 9 and the PWR cycle, described on page 11.

The MFX utility token stores value on the network, facilitates governance, ensures proper network security, and rewards participants to provide work to the network. MFX must be burnt to create PWR, which is used to drive the network. Therefore, generally speaking, MFX is burned with network use, and MFX is minted when nodes supply work to the network.

PWR is a digital credit that is redeemable for network services. It is created via burning MFX. It exists within the network, as it can only be used to redeem services on the network. When used for services, PWR is burnt. Producers of the services are rewarded in MFX. PWR cannot be burned to create MFX. Once a user holds PWR, it can only be used to redeem resources on the network.

While the market price of MFX may fluctuate, and therefore the MFX to PWR redemption rate may fluctuate, the number of PWR needed to redeem specific network resources is set by the governance system. This establishes a consistent cost structure for network resources.

Generally speaking, demand for the network results in more burning of MFX, while supplying services to the network via validating or running an operator node results in newly minted MFX. In other words, demand deflates the MFX economy supply, while supply inflates.

The governance system, discussed in Section [x], controls this burn/mint equilibrium to ensure the economy doesn't hyperinflate due to over-liquidity or stagnate due to under-liquidity.

4.1 MFX Cycle

The MFX utility token stores value on the network, facilitates governance, ensures proper network security, and rewards participants to provide work to the network.

MFX must be burnt to create PWR, which is used to drive the network. Therefore, generally speaking, MFX is burned with network use, and MFX is minted when nodes supply work to the network.

The MFX cycle has inflationary characteristics and deflationary characteristics, which are described in detail below.

4.1.1 Inflationary Drivers of MFX

Staking by Network Operations / Infrastructure Suppliers Suppliers of network fundamentals are rewarded in newly-minted tokens. This includes suppliers of CPU, memory, transfer, bandwidth, and disk space, amongst others. Network operators are further discussed in Section [x].

Governance Activity The governance of the Manifest Network is powered by the MFX token. To incentivize engagement and ensure proper conduct of the governance system, various types of governance activities may be rewarded by newly minted tokens.

Examples could include proposals where the winners are awarded in newly minted tokens, and losers have their tokens slashed. In some cases, submitting a proposal may result in an award of newly minted MFX. Governance activity is described in section [x].

Staking for Security and Liquidity Control The Manifest Network is not a staking yield network. However, in the lifetime of the network, it may temporarily provide yield to stakers for security purposes.

Most yield-based staking protocols have collapsed. Rather than locking up liquidity on the network and generating a high demand of investors to chase yield, the Governance Network may choose staking and yield farming as a utility to control the growth of the protocol's economy by locking up network liquidity.

The Manifest Network will never have people stake for yield to drive up the token price. If the network needs to balance out excessive liquidity, the governance system may elect to reward stakers for liquidity controls. In this case, providers would lock up quantities of tokens, and earn a reward in newly minted tokens, at a rate determined by the governance system. In this case, yields would be inversely proportional to the number of people staking on the network.

To make the risk of yield-farming commensurate with the yield itself, there will always be slashing associated with the stake and its yield. In the upside scenario, these stakers will get a yield, but in a downside scenario, their stake will be slashed and they will lose all their tokens.

This form of slashing can result in a price stabilization mechanism for MFX. When most protocols enter a downside spiral, the only outlet is a decline in token price. Instead, when yields are slashed in a downside scenario, a contraction in the token supply becomes the compensating control when the protocol becomes unbalanced.

In this case, all stakers who are earning yields on non-utility functions will be slashed and all their MFX will be burned.

4.1.2 Deflationary Drivers of MFX

Usage MFX must be burnt to receive PWR credits, which are required to use the network. Therefore, network usage results in a decline of circulating MFX supply. A high degree of demand on the network results in a high volume of burnt MFX, balanced by a high volume of newly minted MFX to reward the Network Operators and Validators that meet the demand.

Transaction fees MFX transactions require a transaction fee. This fee is burnt upon transacting. Transaction fees are set by the governance network.

Slashing When an operator or validator node performs an operation, the nodes' stake is frozen for 30 days. If a node misbehaves or acts in a way that is not to the benefit of the network, its staked tokens can be slashed and their tokens are slashed.

Node misbehavior results in a slashing schedule based on what they did, and a different amount of tokens will be slashed according to whether they are acting in a malicious way, or simply down too much, or just misconfigured.

Slashing in this way prevents bad actors from performing an operation and then leaving the system. All slashing rules are determined through a governance proposal.

Governance Some governance activity results in burning tokens. For example, there may be votes where the winners are rewarded with newly minted tokens, but the losers' tokens are slashed.

Dead loss There will always be a certain amount of tokens for which people lose their credentials, and, since the assets exist in a purely virtual space, cannot be recovered again. Private key holders may die without handing off their keys, or holders may lose their private keys, or assets may get trapped in a corporation when everyone that knows about the assets leave. Dead loss results in some depreciation of the network.

4.2 PWR Cycle

PWR is a digital credit that is redeemable for network services. It is created via burning MFX. It exists within the network, as it can only be used to redeem services on the network. When used for services, PWR is burnt. PWR cannot be burned to create MFX. Once a user holds PWR, it can only be used to redeem resources on the network. An application run on the network must be "charged" with PWR Credits.

PWR Credits are cryptographic tokens only in the sense that they are cryptographically-secured mathematical representations of the ability of the holder to use the network. However, PWR credits do not resemble currency or traditional crypto-economic coins or tokens.

PWR Credits only exist inside the network. A developer building on the network may get PWR by burning MFX and receiving PWR according to the real-time redemption rate. PWR Credits can also be sent between accounts on the network. In this case, a developer could build on the network without ever having to purchase and burn MFX.

The use of PWR Credits on the network ensures that the cost to use the network does not fluctuate with the price of MFX. PWR is not a stable coin. However, because PWR represents the ability of the holder to receive a certain amount of work from the network, as defined in the cost tables in ([coming date]), and the cost of the network is determined by the production cost of network operators, the price is not PWR and is not volatile.

The redemption rate of MFX to PWR is variable. The governance system adjusts the redemption rate to ensure proper network incentives. For example, a user requiring 20 CPU cores, 10 GB of memory, and 10 GB of start disk may expect the cost to be the equivalent of $891.48 of PWR Credits a month. A full breakdown of an example cost table can be found on page nineteen.

4.2.1 Pre-Loading Accounts

The Manifest Network allows the end-user to interact and use systems built on the Manifest Network without being required to hold PWR themselves. This effectively allows the network to resemble traditional web hosting.

Application developers pre-load the cost of execution in PWR on accounts linked to infrastructure that can be used. End-users charge down pre-loaded PWR as they use the application. If PWR is fully spent down, the contract is no longer able to be used, and whatever the contract is powering stops. Cost is pre-loaded by the application developer (the one initiating the contract) and spent by the user (the one utilizing the contract), so the user doesn't actually need to pay when they use the contract. This model allows the end-user experience to resemble Web 2.0. An end-user may make use of services powered by the network without ever buying, holding, or owning MFX or PWR. This reduces the barriers to the adoption of the network.

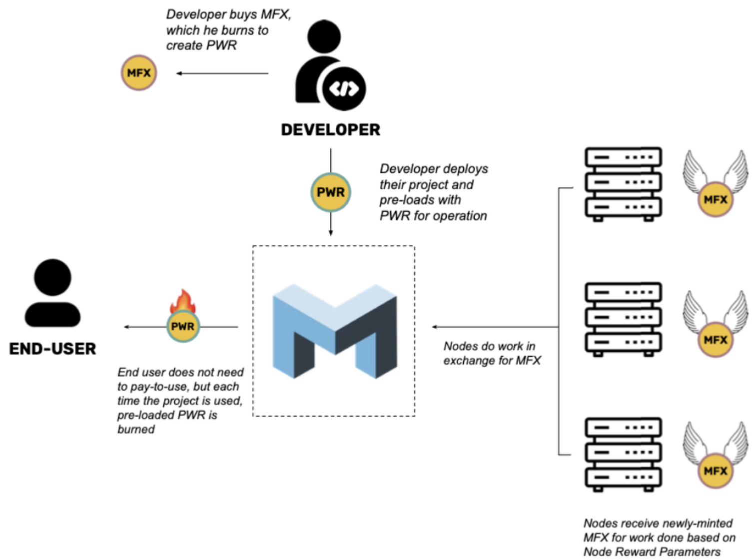

Figure. 1 Manifest Network's Dual Asset Model

4.2.2 PWR Stations

PWR stations are independent operators that acquire and hold MFX to convert into PWR, and thereby allow a user or developer to acquire PWR credits directly without burning MFX. They may charge the user in fiat or any other crypto, and in return they credit PWR to the purchaser's account.

PWR stations are applications external to the Manifest Network, but provide a crucial service to the network. PWR Stations reduce the steps for a user to acquire PWR and use the network without buying and burning MFX. Ultimately, these operators will become the primary method by which a user acquires PWR, and creates a straightforward, user-friendly way for a developer to onboard.

In exchange, power stations may charge a fee for the convenience afforded the user.

Any person or business that acquires MFX may become a PWR station. PWR stations do not have to run a node. There are no conditions or restrictions by the network for an operator to open a PWR station. They simply need to acquire MFX, set up a payment gateway, structure it appropriately, and sell PWR to purchasers.

5 Nodes

Most networks conflate validation, which is a security function, with governance, which is a governance function, with supply of transactions, which is a CPU, memory, bandwidth and storage function. The Manifest Network breaks this all apart into separate networks with discrete functions that contribute to the MFX token economy.

Since nodes are rewarded with newly minted tokens, and thereby inflate the network, the number of nodes authorized by the network is discrete and planned by the governance network. Refer to Capacity Planning in Section [x].

This section expounds upon the tokenomic aspects of the Operational Network, the Validator Network, and the Governance network.

5.0.1 Capacity Planning

Traditional networks tend to be either completely open or overly controlled and centralized. In either case, capacity planning for nodes on those networks tends to be non-existent.

Completely open networks, such as Bitcoin Proof-of-Work or Ethereum Proof-of-Work, allow anyone to become a node. In theory, the cost fluctuations of the network are meant to control the amount of capacity on the network.

However, letting the network self-regulate has increased the cost of transactions for users and constricted network throughput.

In the case of the Bitcoin network, as the price of BTC increases, the value of the mining reward increases, which incentivizes more and more miners to join the network and compete to get the reward. The increased competition of miners, however, causes the total cost of organizing transactions on the network to increase. The competition to be a node is out of proportion to the actual amount of work needed to supply a transaction, which increases cost of that transaction and constricts network throughput.

Consequently, open node networks that token price to self-regulate capacity have not proven themselves effective for production-oriented networks that require a high volume of throughput.

Conversely, a rising cadre of new chains including Solana, Ripple, BNB, tightly control their networks. These chains avoid the cost and throughput problem of open networks at the cost of highly centralizing their network.

The Manifest Networks' governance system regulates node capacity in proportion to network demand and consumption, without excessively centralizing the network. The governance system will moderate the growth of the network over time, and expand the capacity of the network based on trends in demand. It will monitor the operational demand of the network and authorize an increase in nodes to meet that demand.

Initially, proposals to expand the capacity of the network will happen by individual voters. Ultimately, capacity planning on the network will be largely automated. The governance network can leverage machine learning nodes to automatically generate proposals to expand or contract the network. However, there must always be human-run voting in any proposals to make sure that the system can't easily be gamed.

5.1 Operator Nodes

Operator Nodes provide the CPU, Disc Space, Bandwidth, Memory, etc. that is the basis for the computational fabric that drives the network. These nodes are rewarded a base amount of newly-minted MFX for general availability to the network, and they are rewarded an additional variable amount of newly-minted MFX when the nodes' resources are used, based on a number of specific performance and production criteria. Since the addition of operator nodes to the network inflates the MFX supply, the number of operator nodes on the network are discrete and defined by the governance system. To become a node and earn MFX, node providers must offer their nodes to the network, which tests and validates the node's resources before inducting the node into the network. Nodes earn base rewards for their resources being available on the network, and variable rewards when those resources are actually used.

5.1.1 Offering a Node to the Network

An increase in the number of nodes on the network is authorized by the governance network via governance proposals. The governance network, which monitors CPU, Memory, Disc Space, trends in demand for these, location of demand, etc., will evaluate when new capacity is needed on the network. If new capacity is needed, the governance system will authorize a new node in a process detailed in Capacity Planning, Section [x].

When a new node wishes to join the network, it makes an offer to the governance network and specifies the resources it can provide and its location (i.e. 400 CPUs in Chicago). The operator signs their offer with their private key so the network can associate the node's resources with their account.

The network then runs a series of tests on the newly-offered resources to validate their capacity. If these resources are appropriately validated, the network will accept the nodes' offer and it will join a queue of other nodes whose resources have been tested and accepted.

When network capacity must increase, the governance system inducts a node from the queue to expand the capacity of the network based on the specific resources and location of that node. Inducted nodes are assigned a base reward for general ability. When they begin to do work in the system, they are rewarded with variable rewards based on specific criteria.

5.1.2 Base Rewards for Resource Availability

Base rewards are provided for the node based on the nodes' availability to provide resources to the network.

Base rewards derive from the number of workloads and/or computations a node can provide. Measurements include, but are not limited to: CPU, Memory, Transfer, Bandwidth, and Disk Space.

Base awards are measured over time. Nodes are only rewarded only base rewards for the time that the node is active. The node does not receive rewards for any time that it is offline.

5.1.3 Variable Node Rewards

In addition to base rewards for availability, nodes are awarded newly minted MFX as their resources are used for the type, quality of, and extent of work. As nodes do more work, they are rewarded more newly-minted MFX.

Variable node rewards are based on Network Reputation, Work Provided, Environmental Offset, Overall Network Performance, Geographical and other offsets.

Network Reputation A node's reputation is variable based on how it behaves. Its reputation is driven by the amount staked in the node, the nodes' vintage, the amount of time the node has continually been online, and the quality of decisions that node has historically made, amongst other factors (to be detailed in future WP).

Stake The more MFX is staked in the node, the more the node loses if it misbehaves, and therefore the higher reputation that node has.

Persistence The earlier the node joins the network, the higher its reputation.

Age The longer the node has been continuously online, the more available they are, and the higher their reward.

Quality of Decisions If the node is frequently out of sync in consensus or acts divergently from the majority of nodes, its reputation, and therefore its reward, will go down.

Work Provided to the Network Work provided is measured by disk, bandwidth, CPU, and memory. The more work provided by a node, and the more successful computation that's done, the more the node will be rewarded. These 4 types of work are measured discretely because they have separate production costs. The amount of MFX rewarded to a node per unit of resource provided to the network will be determined by the governance network.

Environmental Offset Nodes will receive MFX rewards for their environmental offset.

An Environmental Foundation will function as a rating agency that evaluates the environmental offset of nodes and measures the overall sustainability of the network. Nodes will submit cryptographically verified environmental ratings along various axes to report their climate impact and, based on those ratings, will be awarded newly minted MFX.

Nodes that operate in a climate-neutral way will receive a base reward. Nodes that act in an environmentally positive way will receive the base award plus a surplus. Nodes that operate poorly or do not submit ratings will be considered in the bottom category, and will not receive any reward.

When a node receives less than the base award amount, the differential (base - reward) will be granted to the environmental foundation. The foundation will then offset the environmental cost of the node to ensure the network is environmentally neutral.

For example, Google Cloud Platform (GCP) is already climate neutral, so if the node is run on GCP it will be able to sign an environmental attestation that declares the node to be climate neutral. The Environmental Foundation will be responsible for verifying that attestation and, once verified, the node will receive the base MFX environmental reward.

The environmental ratings system will be described in a future whitepaper.

Overall Network Performance and Goal Achievements The incentives of node providers need to be aligned with the overall health of the network. The governance network, therefore, will be capable of establishing entire network performance targets. Upon successful completion of these targets, all node providers will receive MFX bonuses.

For example, the governance network may provide incentives to expand the network's geographic distribution, or increase the number of small nodes operating on the network, or any other specific goals that benefit the entire network to incentivize collective action.

Network Performance Rewards help to encourage specific behavior that contributes to the overall value of the network.

Geographical and Other Specialty Rewards Some geographic locations will struggle to be serviced without proper incentive. For example, there has been a persistent historical lack of network infrastructure in Africa. The governance system may create specific rewards for geographical offset, where nodes stood up in underrepresented regions are rewarded extra.

Additionally, other specialty rewards may be assigned to data centers that demonstrate increased resilience. For example, nodes that are nuclear resistant or that have an EMP shield may receive an operating bonus.

Any governance specialty rewards will have a default lifespan of 12 months, and therefore must be renewed to continue.

5.2 Validator Nodes

As with operational nodes, the number of validator nodes accepted in the network will be discrete to prevent overinflation.

The validator network will operate independently of the operational node network. The validator network will look at the transactions happening everywhere on all the chains, and check the cryptography of all of those transactions.

5.2.1 Validator Network Caching Function

The validator network may also cache all of the data from a network. This provides an entire-network cache.

For example, some execution networks might be programmed to delete their data every day. Because the history is not needed, only the latest hashes are needed from a security standpoint.

Validators may choose to provide caching-as-a-service and store past data from these networks, and may monetize this data in the future.

5.3 Governor Nodes

The governance network is responsible for the governance of the entire Manifest Network token economy. The governance network will have its own nodes that do the work of voting, organizing and generating proposals, gathering data required to generate proposals, etc.

All governance activity is conducted using MFX, and governor nodes may gain or lose staked MFX based on various governance functions. Some votes might require staking, where nodes must stake against the success or failure of the proposal, and the winners take the whole pie.

The details of the governance network will be expounded upon in a future white paper. Some areas where governance activity interfaces with the manifest token economy are listed below.

5.3.1 Governance Function in the Token Economy

Burn/Mint Equilibrium (BME) The governance function is responsible for balancing our inflationary characteristics to maintain inflation at an acceptable level. If inflation gets too high, the governance function will be responsible for increasing the incentive for deflationary behavior on the network (e.g. adjusting the MFX / PWR real-time redemption rate to incentivize more burning of MFX). The network liquidity expands and contracts MFX liquidity on the network based on actual operational usage. This process of BME is overseen by the governance network.

Token Redemption Rates The governance system determines how much PWR is received for burning MFX.

Cost Tables The governance system is responsible for determining how much PWR is charged for computation, storage, transfer, etc.

Example cost table:

CPU	**20**	vCPU
Memory	**10**	GB
Start Disk	**10**	GB
Monthly Cost	**891.475**	

Cores \ Memory	1	2	3	4	5	6	7	8	9	10
2	89.125	94.725	100.325	105.925	111.525	117.125	122.725	128.325	133.925	139.525
4		178.3	183.9	189.475	195.075	200.675	206.275	211.875	217.475	223.075
6			267.45	273.025	278.625	284.225	289.825	295.425	301.025	306.625
8				356.575	362.175	367.775	373.375	378.975	384.575	390.175
10					445.725	451.325	456.925	462.525	468.125	473.725
12						534.875	540.475	546.075	551.675	557.275
14							624.025	629.625	635.225	640.825
16								713.175	718.775	724.375
18									802.325	807.925
20										891.475

Node Rewards The governance system is responsible for determining how much nodes earn based on the node parameters described in this paper. The governance system may also issue new reward parameters.

Node Capacity Planning The governance network is responsible for authorizing new nodes onto the network. The amount of PWR burnt by network users signals to the governance network that more demand is coming into the network, which is a signal to increase the operational capacity of the network. The governor network then authorizes more nodes, which are drawn from the cue of available operation nodes. The authorization of more nodes increases the MFX liquidity within the network via newly minted node rewards.

The governance system must also approve a node to leave the system, to ensure the node hasn't performed as a bad actor.

6 Appendix: Dapp User Experience

The experience that a Dapp user has is up to the Dapp developer. They have at least 4 options they may choose from to provide an optimal experience for their end-user. The use cases for the end-user experience allow for full compatibility with both web2 (centralized) and web3 (decentralized) projects. A fully centralized project could use the Manifest Network without its users experiencing any friction with onboarding. A fully decentralized platform, likewise, can also build on the Manifest Network, without relying on a centralized service provider.

6.1 Free

A Dapp developer may choose to not charge the Dapp user anything. This may be the case of a free system, such as a free social network, or a Dapp that collects data and monetizes that data on the back end, a charity, or any other uses a Dapp developer can come up with.

In this case, the Dapp developer pre-pays for usage with PWR and simply lets the user use the system.

6.2 Charge the End-User through Traditional Payment Pathways

A developer can also use the network in nearly an identical way as current web2 systems. They can add a payment gateway or a regular storefront to charge the end-user. The end-user, in this case, may not even be aware that the application is decentralized.

In this case, the developer would receive money as revenue from their users, in the exact same way they do now, and as their variable costs to host their system on our network would be expensed in MFX converted to PWR, or directly buying pwr from a power station.

6.3 Direct Tokenized Experience with MFX

If the developer is building a fully decentralized application, they may not be interested in revenue from the end-user, or may only want to transact with tokens.

In this case, the Dapp would require the end-user to actually hold and use MFX, pay with MFX on the Dapp, some of which would be converted and used as PWR and some of which would be collected as a fee for the Dapps service (profit).

6.4 Dapp with Its Own Token

Some applications building on the network will want their own tokenized economy, with their own token. Since the Manifest Network is a multi-token network, they can create their own token on the network. They may use their own network (unique consensus algo etc.) built with our tools that federate with the many protocols and become part of the Manyverse.

In this case, their economy would be independent of the Manifest Network economy. However, fees would be paid in MFX and the utility of the network would be paid in PWR.

7 Acknowledgements

In this section we thank the members of the academic team at the University of Texas at Austin (Sandeep Chincali, Denzell Ford, Federico Torres, and Thad Hutcheson) for their help in providing code and support of code for tokenomics simulations, and review and suggestions for edits to this white paper. We further thank Felix Morency for the introduction to the Overleaf tool and contribution to the literature review.

References

[1] G. Adams. *Burn and mint equilibrium.* 2019. URL: https://messari.io/report/burn-and-mint-equilibrium.

[2] O. Akcin et al. *A control theoretic approach to infrastructure-centric blockchain tokenomics.* URL: https://doi.org/10.48550/arXiv.2210.12881.

[3] V. Buterin. *Ethereum Whitepaper: A next-generation smart contract and decentralized application platform.* URL: https://ethereum.org/en/whitepaper/.

[4] A. Donmez and A. Karaivanov. "Transaction fee economics in the Ethereum blockchain." In: *Economic Inquiry* 60 (1), pp. 265–292.

[5] J. Kwon and E. Buchman. *Cosmos Whitepaper: A network of distributed ledgers.* URL: https://v1.cosmos.network/resources/whitepaper.

[6] S. Nakamoto. *Bitcoin: A peer-to-peer electronic cash system.* URL: https://bitcoin.org/bitcoin.pdf.

[7] T. Rocket et al. *Scalable and probabilistic leaderless BFT consensus through metastability.* 2019. URL: https://whitepaper.io/document/604/avalanche-whitepaper.

[8] G Wood. *Polkadot: Vision for a heterogeneous multi-chain framework.* URL: https://whitepaper.io/document/596/polkadot-whitepaper.