09 January 2023 Litepaper V1.0

# The Decentralized Computational Network

## Blockchain-as-a-Service solutions built on the blockchain

**Proof of Useful Work - Litepaper**

Written and maintained by The Flux Development Team

Alex Perritaz, Davy Wittock

Sponsor - Daniel Keller

runonflux.io

# Table of Contents

# Introduction

This paper provides a comprehensive description of the Flux Ecosystem, with a particular emphasis on the Proof of Useful Work implementation. The Flux Team aims to make this paper accessible to everyone by providing non-technical summaries of the entire ecosystem. Individual technical papers for each project component are in-process and will be released as they become available.

## Executive Summary

The Flux PoUW technical white paper provides detailed technical information for both types of users: clients and providers. This document discusses the system architecture including the consensus algorithm of the system, hardware description, architecture of the platform, limitations, and tokenomics.

## Disclaimer

This document is provided solely for reference purposes and is subject to change in subsequent versions. The technical whitepaper and GitHub repository will be updated later this year with additional information regarding the architecture of the provisioning service, the development of use-case-specific images, and the tokenomics.

# Proof of Useful Work

The Proof of Useful Work initiative by Flux lays the foundation for a more conscious and sustainable approach to IT resource utilization. By pooling idle resources and focusing on real-world problem-solving as opposed to transaction validation, we established one of the world's largest community-powered computational networks.

Flux platform allows providers and individuals to rent out their hardware for tasks like AI training, content rendering and cloud gaming. These hosts have an option to schedule time when their machine can be borrowed by clients. The platform's interface allows them to monitor, configure and set network preferences.

The platform provides clients who wish to rent specific hardware with a search and matching interface. They are able to search the network's marketplace for options that meet their computational needs and rent additional computing resources to speed up their work.

## Use cases

The current infrastructure provides a reliable, secure, and scalable ecosystem while maximizing decentralization, required for the implementation of a PoUW consensus, which will secure blockchain mining networks while solving rendering, machine learning, deep fake detection, and other computational lifting problems.

### Artificial intelligence

Artificial intelligence describes the simulation or approximation of human intelligence in machines. It encompasses techniques ranging from machine learning to deep learning, which have applications in numerous fields, such as medicine, academic research, and weather forecasting, among others.

Academic research teams, data scientists, and other institutions are outfitted with extensive computing resources to meet specific requirements. Currently, training and computation problems for deep learning models in organizations require a considerable amount of time.

Flux PoUW offers machine learning, cloud computing, and other AI-related services at highly competitive prices and with uninterrupted uptime. Run workloads concurrently until completion with minimal user intervention. In addition, when a container terminates, the released resources can be allocated to other workloads, resulting in increased system overall efficiency.

### Rendering

Rendering is the process by which an application program generates a two-dimensional or three-dimensional image from a model. It is predominantly employed in architectural designs, video games, animated films, simulators, television special effects, and design visualization.

With a decentralized rendering service, creators can enjoy all the advantages of a render farm without its high costs associated with its maintenance, hardware, and software.

### Headless Gaming

Playing video games without the need to install them on your computer or console, and the need to purchase a graphics card of the latest generation to appreciate it fully, is what decentralization contributes to.

Using someone else's idle GPU to play the latest AAA games on your laptop or even in-browser on your 4K smartTV, enabling uses of advanced display technologies such as 4K, VR, and AR allow users to immerse themselves in these games with hyper-realistic graphics and intricate elements in expansive virtual worlds without having to purchase a new GPU with each new release.

## Digital Asset

Through network participation and resource sharing, Flux provides an innovative, efficient and flexible method of using time and resources available.

Flux utility tokens can be used as a computational battery, accumulating over time through hosting nodes for resource sharing or providing useful work, rewarding contributors based on the quality of service provided.

The accumulated asset can also be used when necessary to rent other services on the platform or to perform and accelerate intensive computations by utilizing and renting additional resources from the network.

## Environmental & Economic impact

PoW has some form of a gambling mechanism: either you discover a block and receive a substantial reward, or you receive nothing. Transaction finality is probabilistic, not guaranteed. After a transaction is included in a mined block, there is always a chance that a longer chain comes along and reverses the transaction, although this probability decreases as additional blocks are added. PoUW, in contrast to PoW, selects the best provider who meets the service requirements, so there is no needless rush to claim a block reward. The tasks are distributed to providers based on their respective hardware specifications.

The lottery effect of miner rewards leads to high payout variance, incentivizing miners to consolidate or form centrally directed pools. This centralization problem has a significant impact on the environment. According to the International Energy Agency, data centers consume approximately 200 terawatt-hours (TWh) of electricity, or nearly 1% of global electricity demand, contributing to 0.3% of all global $CO_2$ emissions. Moreover, concentrating processing power in a single data center location is expensive. Owners have to cater to the system’s risks about security and resource availability, administration, and maintenance. PoUW reduces the energy overhead of mining by decentralizing processing power.

By requiring a large amount of energy per transaction in PoW, the fee per transaction (either direct or indirect via inflation) must be high enough to make up for that real cost. Lowering the energy per transaction of course, commensurately lowers the security of each transaction. Flux is presently implementing the necessary infrastructure to support the PoUW consensus, which seeks to secure transactions based on computational tasks.

### Efficient use of resources

Given that the majority of the hardware is idle or operating at its lowest capacity for the majority of the time, this can be considered an inefficient and wasteful use of these resources. The investment in these is commonly justified for a practical need of that resource, but it is mostly unjustified compared to the resource's actual usage.

By pooling resources and computational power, storing it over time as an asset in Flux, and using what we need, when we need it, we can extend the lifespan of hardware components and limit obsolescence, thereby reducing the frequency of hardware refreshment and maximizing the use of existing resources.

## Conformity with United Nations objectives

The project makes it possible to establish connections with UN Sustainable Development Goals 9 (Industry, Innovation, and Infrastructure) and 12 (Climate Action) (Responsible Consumption and Production). The 2030 Agenda permits these goals to be broken down into sub-goals, and the proposed project has a direct relationship with several of these sub-goals.

Specifically, Goal 9 aims to "promote sustainable industrialization that benefits all." Pooling makes available to the community any surplus resources. It assists in rationalizing IT resource utilization. This aligns with the 9.4th objective of the 2030 Agenda: "Modernize infrastructure and adapt industries to make them sustainable by 2030 through more efficient use of resources and increased use of environmentally friendly technologies and industrial processes, with each country acting within its means.

Objective 12 seeks to establish consumption and production patterns that are sustainable. Resource efficiency contributes to the reduction of obsolescence and, consequently, to the reduction of waste generation, as stated in Objective 12.5: "Through prevention, reduction, recycling, and reuse, significantly reduce waste generation by 2030.

# Ecosystem

The current network infrastructure consists of FluxNodes, a distributed incentive-based hardware solution that can accommodate multiple modules. These modules include FluxCloud, a decentralized storage solution; FluxGBPs, a network traffic aggregator that enables high-bandwidth alternatives for large data transfers; and proof of useful work (PoUW), a decentralized computing solution.

| Module | FluxNodes | PoUW | FluxCloud | Flux GBPs |
|---|---|---|---|---|
| Main Function | Infrastructure | Computation | Storage | Bandwidth |
| Second Function | dApp | Security | N/A | N/A |

The Proof of Work (PoW) component will be added to the Proof of Useful Work (PoUW) package, ensuring that the network is secured correctly and that a reasonable block time is maintained during the transition.

| Module | PoW | PoUW | FluxNodes | FluxCloud | FluxGBPs |
|---|---|---|---|---|---|
| Network Security | 100% → 0% | 0% → 100% | 0% | N/A | N/A |

This balance between the PoW and PoUW consensus will be dynamically adjusted and will decrease as adoption grows and workload increases.

## Proof of Useful Work

The components tied to the PoUW provisioning service consist of an image library for each use-case and the most common field of application-specific configurations, a connection to a specified storage solution, all defined by the specifications made by the client through a use-case specific marketplace.

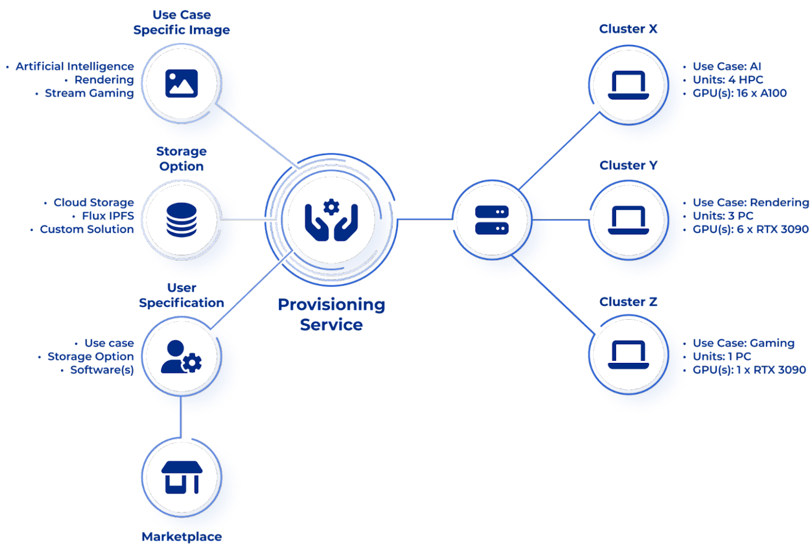


The provisioning service is responsible for the delivery of each image and configuration to the selected machines. It provides information to a management interface for the cluster's worker nodes.

# Utilization

The hosts download and install the management software, then list their machines, configure them, and pre-configure for any default jobs.

When configuring the management software for PoUW, hosts can specify how and when they want their resources to be shared to the network;  they can either leave it on PoUW full time while securing the network when not assigned to a task, be on standby and switched on only when a task is available, or be put on standby and wait for a task at a defined schedule.

Clients can then access these through the marketplace, giving them access to a vast array of machines and configurations and the ability to rent suitable ones and execute the desired task within a few clicks.

## Hardware

Most components can contribute to PoUW in some capacity, but the ability to run specific use-cases will be restricted to a defined list of hardware that meets the minimum requirements for hosting and running these services at their lowest settings.

These requirements will serve as the foundation for the current iteration, but they are subject to modification to define the minimal requirements while ensuring computation quality.

### CPU & Motherboard

Hardware virtualization and IOMMU must be supported by the CPU for GPU passthrough. IOMMU must be supported by both the motherboard's chipset and BIOS[1].

| Intel | AMD |
|---|---|
| Pentium 4 / Core I3, I5, And I7 / Xeon | Bulldozer generation and up (including Zen) |

### Bandwidth

The minimum bandwidth and latency requirements are defined by each use-case.

- 4K Gaming : 50-100 Mbps download / upload
- Regular gaming : 10-25 Mbps
- File transfer : 5-10 Mbps download / upload

[1] https://handwiki.org/wiki/List_of_IOMMU-supporting_hardware

### GPU

From Nvidia the following components :

- Turing (RTX 20 Series, Titan RTX, GeForce RTX, GTX 16 Series)
- Pascal (GTX 10xx)
- Maxwell (GTX 7xx, GTX8xx, GTX9xx)
- Kepler (GTX 680, GTX 690)
- GTX Titan
- Quadro GPUs
- Volta GPUs

At the time of writing, AMD and Intel GPUs support no guarantee of their usability for all use cases.

## Benchmarking

When first installed, a preliminary evaluation of the hardware is performed to ensure the service's quality, to assess the performance for each use-case, and to accurately list the machine and its components on the marketplace. The execution of an additional benchmark precedes every new job to ensure that the configuration specified is valid and can perform as expected.

The components benchmarking score is computed by a score defined by the following elements:

- GPU
- CPU
- Memory
- Storage
- Bus speed
- Network speed / Latency
- Max. Availability

In order to provide up-to-date data and ensure service quality over time, random benchmarks are performed frequently throughout the lifetime of these machines on the network.

**Note** : Each benchmark provides a score that is unique to the assigned use case.

## Storage

PoUW heavily relies on the storage component to ensure that all systems are running tasks and limit the amount of idle or unused resources. Consequently, it is essential to release systems that have completed their assigned tasks by uploading data, intermediate results, and final results to a persistent location.

The PoUW infrastructure enables users to specify additional work directories and storage locations, thereby making it S3 compliant. This option facilitates the adoption of the decentralized computation infrastructure and integration of the computation component into a more conventional workflow.

# Technology

The key components of Proof of Useful Work implementation relies heavily on the current node infrastructure, which ensures network decentralization and the hosting of various services, Kubernetes for task orchestration and delivery across the network, and finally the PoUW consensus, which will be incorporated within the finalized architecture.

## Architecture

At its core, PoUW infrastructure heavily relies on Kubernetes to facilitate the orchestration of the operations required for the automatic provisioning of cloud services, including the creation of the virtual IT infrastructure, the installation of necessary software components, their configuration, and ensures proper execution of tasks.

Building the architecture of the network around a lightweight flavor of Kubernetes, with a special focus on scalability and adaptability to any type of cloud environment and use cases, with built-in integration for any existing ecosystem.

The abstraction of GPUs enables efficient sharing and automated configuration of these resources across multiple workloads for each use-case while ensuring dynamic allocation of GPUs to workloads that require them, optimizing their usage and eliminating the need for pre-allocation of future capacity needs.

### Artificial Intelligence

Kubernetes and containers are central to PoUW, and AI use cases due to their capacity to decouple modern distributed applications from the infrastructure layer.

Access GPU components via Kubernetes device plugins and GPU-enabled dockers to leverage high-performance, scalable infrastructures operated by individuals and utilizing pooled resources to power AI workloads.

#### Specific containerized solutions

Start training your models by deploying your favorite AI frameworks and libraries directly within the containers. Construct, train, and deploy models for any use case and accelerate adoption by enabling the integration and deployment of any containerized solution with GPU-ready instances on the platform.

### Rendering

Depending on the software and field of application of the desired task, preconfigured images with the following options will be available to deploy onto the selected machines. The following options may change over time to include new or alternative tools.

#### Option 1 : Octane Network Render [2]

"Octane Network rendering is a GPU-based rendering that allows additional computers to be utilized for image rendering. OctaneRender distributes compiled render data rather

[2] https://www.cadnetwork.de/documents/Guide_Network_Rendering_With_Octane.pdf

than scene data, eliminating the need for the user to manage files. For network rendering, a master computer and one or more worker computers are required.”

### Option 2 : Pulze [3]

Pulze is an artist-friendly solution that integrates the three solutions that will accompany the artist throughout the project's lifecycle. Scene Manage, Render Manager, and Post Manager are the tools that allow you to collect all of your 3ds Max settings and bring them into a unified interface that is simple to use and allows you to leverage rendering on multiple machines for efficient and easy management of all stages of content creation. Pulze currently supports the most commonly adopted 3ds Max software.

## Headless gaming

We’re basing our headless gaming images on the following GitHub projects, which allows anyone with a modern web browser, to start playing their favorite AAA title games on any type of device while also supporting in-game controller integrations.

### Option 1 : Docker Nvidia GLX Desktop [4]

“A container created for Kubernetes that supports OpenGL GLX and Vulkan for NVIDIA GPUs with WebRTC and HTML5, offering an open source remote cloud graphics or game streaming platform. Instead of using the host's X server, it spawns its own isolated instance.

### Option 2 : Sunshine & Moonlight [5]

“Sunshine is a Game stream host for Moonlight. Sunshine is a self-hosted, low latency, cloud gaming solution with support for AMD, Intel, and Nvidia gpus. It is an open-source implementation of NVIDIA's GameStream, as used by the NVIDIA Shield. Connect to Sunshine from any Moonlight client, available for nearly any device imaginable.”

> **Note :** Options discussed in this section are known to operate at larger scales and are used exclusively as proof of concept and to demonstrate the viability of the approach.

[3] https://www.pulze.io/
[4] https://github.com/ehfd/docker-nvidia-glx-desktop
[5] https://github.com/LizardByte/Sunshine

# Privacy

Our privacy policy regarding the collection, handling, and processing of hosts' and customers' data will be kept strictly confidential and will not be sold or distributed to third parties.

## Sensitive Data

A special emphasis is placed on ensuring that individuals' and organizations' confidential information that requires protection from unauthorized disclosure remains out of reach to unauthorized parties. Without the prior approval of the user, data should never be moved or transferred to another computing unit.

## Secure enclave

As hosts are not governed by a legal entity, we aim to provide a secure environment restricting their access to areas of their shared resource that may contain sensitive data. Providing CPU hardware-level isolation and memory encryption on every server by isolating application code and data from any privileged user, with additional storage and network encryption on top.

## Scope of use

The geolocation of each host and client can be specified to provide services that adhere to specific country policies and GDPR compliance requirements. The same is true for organizational and institutional entities, which can only be specified and shared within the scope defined.

# Life Cycle

## Overview

Desired State Description Language default configuration specification :

- Installation of PoUW core image
- Run Benchmarking - general purpose
- Run Benchmarking - use-case specific
- Link Storage information

## Installation on the machine (Host)

**[Software/OS installation]**

1. Installation on local machine
    a. Link to wallet
    b. Register/Login with host UID
    c. Confirmation / Activation
2. Run general purpose benchmarking
3. Run all use-case-specific benchmarks
4. Send benchmark scores, UID, and wallet ID to logical database
   *Logical database : price calculation, price attribution, …*
5. Define availability
6. Wait for work

## Perform Task (Host machine)

**[Job received]**

1. Save the initial state locally
2. Use-case provisioning
3. Use-case benchmarking [fail | success – % threshold]
4. Load user data (defined in package specification)
5. Perform task
6. End task
7. Upload user result (defined in package specification)
8. Update machine performance
9. Reset the initial state locally
10. Dismount

# Website

## Marketplace

The marketplace is divided into sections devoted to each use-case and provides users with popular pre-configured images for the most prevalent field of applications for that use-case, which can be deployed within a single click on all the selected machines.

Users also have the option to configure and tailor these to their specific needs, or the ability to start with an empty configuration on which they specify all required components from scratch.

### Artificial Intelligence

**Home FluxAI > AI Ready**
Field of Application > Refine selection > Select preset > Select additional components (opt) > Display available machine(s) > Select machines > Deploy

**Home FluxAI > AI Custom Selection**
Display available machines > Refine selection (opt) > Select preset > Select additional components (opt) > Select machine(s) > Deploy

### Rendering

**Home FluxAI > Render Ready**
Field of Application > Refine selection > Select preset > Select additional components (opt) > Display available machine(s) > Select machines > Deploy

**Home FluxAI > Render Custom Selection**
Display available machines > Refine selection (opt) > Select preset > Select additional components (opt) > Select machine(s) > Deploy

### Gaming

**Home FluxGaming > Game Ready**

Home FluxGaming > Desired Settings (4k, 1080p, ...) > Games available > Filter category (opt) > Select game > Rig available > Select machine > Deploy

# Client UI

**Dashboard Features**

- *Billing* : Wallet / Funds / Costs / History
- *Rented resources* : Configurations / Usage / Costs / History
    - Nodes resources : Available / Usage / Allocated
    - Component metrics : Idle / Usage / Allocation
    - Jobs : Running / Usage / Allocation / Pending
    - Additional info. : Geolocation / Organization
    - Shared account : Jobs per user/team

**Data repository**

- Flux IPFS
- Cloud provider
- Git repository hosting services
- Custom solution

**Projects repository**

- Resource allocation (per project)
- Set limit : Time / Cost / Usage / Allocation
- Metrics : Usage / Cost / Assigned

**Market Place**

- *Select* pricing model (On-demand, Pay-as-you-go, Spot market)
- *Display* all machines
- *Select* hardware configurations and use-case
- *Select* machine(s)
- *Select* use-case (AI/Rendering/Gaming/Other)
- *Select* use-case components to be installed (presets or custom)
- *Select* data source
- *Deploy* selected machines (configurations & components selected)

# Host UI

**Dashboard Features**

- *Wallet (Payouts in Flux)*: Funds / Costs / History
- *Machines* : Configurations / Usage / Profits / History
    - Nodes resources : Idle / Usage / Allocated
    - Component metrics : Idle / Usage / Allocation
    - Jobs : Running / Usage / Allocation / Pending / Profits
    - Additional info. : Geolocation / Organization
    - Settings : Availability / Hardware limitations / Selected Use-case

# Tokenomics

## Price structure

The price is composed of three primary components :

- Active rental cost
- Storage rental cost
- Bandwidth utilization cost

The *active cost* is the cost per second that an instance is in an active/connected state, while the *storage cost* is charged for each and every second that an instance exists, regardless of its state. The third component is the *cost of bandwidth usage*, which is based solely on the volume of data transferred over the network.

### Active rental Cost

The pricing of an instance is determined by its configuration and assessment of its components for a given use case, in addition to the subsequent factors :

- Pricing model
- Baseline price
- Supply / Demand
- Reputation score
- Benchmark score by use-case
- Availability
- Verified status

**[Formula & Explanation TBD]**

**[Pricing calculation TBD]**

### Penalty - Interruptions and biased benchmarks

If a job is interrupted before completion, the user's entire payment will not be lost, and the current state will be resumed on a similar or equivalent machine if one is available. If partial results are available, only the completed work will be credited.

Jobs that are interrupted or machines that perform poorly relative to their benchmark score will have their reputation score decreased.

# Asset distribution

The economics of decentralized computing infrastructure and the prices charged to customers for decentralized resources are governed by three pricing models inspired by Amazon but at a fraction of the price. Following are the specifics of the distribution of the flux digital asset after payment for a completed task.

## Client

Depending on the chosen payment method, an additional network fee is punctuated for the use of the decentralized computational services when using a fiat on-ramp option, while the fee is reduced when using a Zelcore wallet. Payments can be made either through fiat or in native Flux.

## Hosting

The entirety of the calculated price is attributed to the host, with the foundation receiving only a nominal transaction fee. Payouts are made to your Zelcore wallet in Flux; fiat payouts are not supported at this time.

## Foundation

The foundation contributes to providing the infrastructure of the decentralized network, ensuring the quality of its services, and reinforcing the network's development through a small transaction fee.

# Acknowledgments

## Partnerships

### Massachusetts Institute of Technology

The Massachusetts Institute of Technology deserves special credit for inspiring and proposing the Proofs of Useful Work consensus, which this network subsequently adopted. [PDF]

### Geneva University of Applied Sciences - Partnership

This endeavor would not have been possible without Geneva University of Applied Sciences. We would like to express my deep gratitude to the Geneva University of Applied Sciences team for their help in offering the resources to make PoUW research.

## Team Members