

FLUENCE

Decentralized database network

July 18, 2017, v1.1

FLUENCE	1
Abstract	2
Fluence	3
Database	3
Open Data Market	4
Use Cases	5
Data Regulation Compliance	5
Data Sharing and Exchange	5
Medical Data Exchange	5
Data Marketplaces	6
Decentralized Apps	7
Machine Learning and AI	8
Economics	8
Fluence Token	8
Business Model	9
Tokensale	10
Budget Allocation	11
Ecosystem Fund	11
Team	12
Roadmap	14

Abstract

We live in a world of big data. Digitalization of life brings convenience in everyday processes and offers many economic benefits. Companies figured out that data insights give market advantages and started aggregate as much data as they can. Currently, pretty much every human's move and step are tracked and stored somewhere in private cloud.

Ubiquitous data tracking brought concerns about sensitive data safety that drove to an emergence of laws and government regulations. However, last few years we have witnessed massive data leaks across many industries. Centralized data storages fall under sieges.

Nobody feels safe in such situation. Modern AI algorithms and technologies demand high amounts of valuable data to solve problems, detect diseases, optimize economics. This requires constant data exchange and aggregation from different sources and automatically enables many threats for data safety. Lots of sensitive data remain in the hands of people and can't be contributed for humanity help.

Blockchain technology gives a new way for storing and handling sensitive data with guaranteed privacy. With modern cryptographic techniques, it is possible to create a decentralized database that allows encrypted storage, data sharing and exchange. Besides creation storage economy, decentralization provides replication of data, permanent retrievability and cost reduction by design.

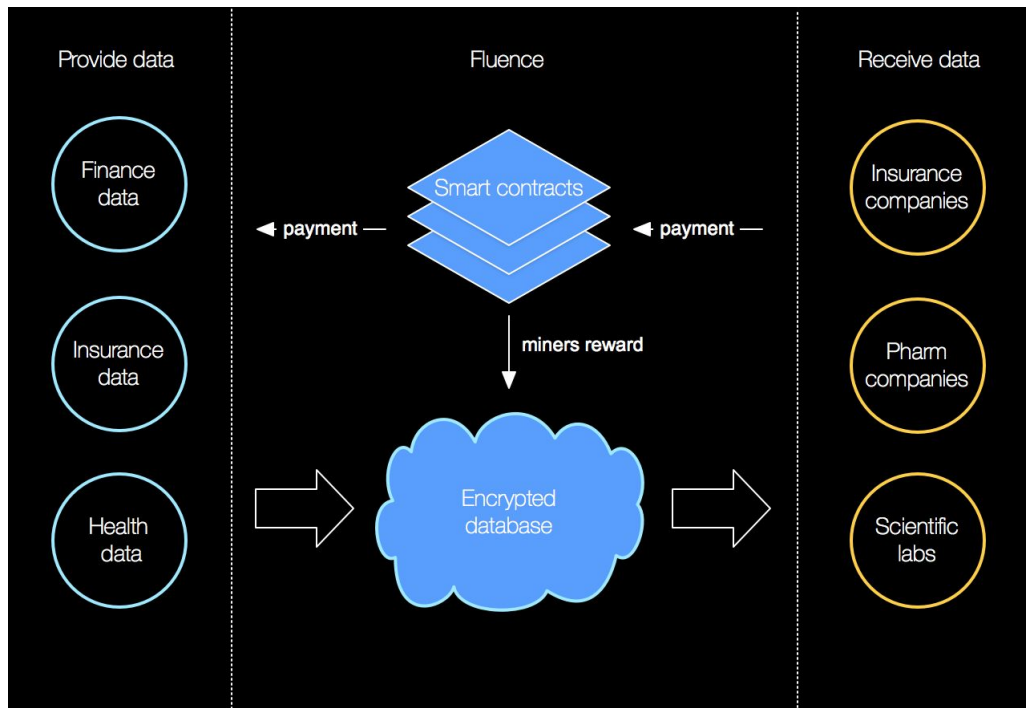
It unlocks many data markets and use cases, where contribution with data is justly rewarded by unbiased smart contract to all providers. For companies and researchers, it breaks barriers of creating value on top of data.

Fluence

Fluence consists of two essential parts: decentralized database and open data market.

To create decentralized, fault-tolerant, censorship resistant database for any structured data, Fluence builds ecosystem when storage owners are incentivized for renting their capacities to data owners.

Data market is represented by low-level API interface on top of the database and centralized website.



Database

Fluence works as a traditional remote database, allows to query data, update or delete with keeping access in secret between you and people you authorized. The major distinction from projects like Storj or Sia is: these projects are basically “decentralized Dropbox”, but Fluence is designed for structured data, not files. Fluence develops API interfaces for app developers to use it as a common database.

In comparison to traditional databases Fluence provides:

- **Fault tolerance and censorship resistance.** Due to decentralized nature, Fluence guarantees maximum uptime and no possibility for data regulation. It automatically balances data replication to keep all data online at any moment.
- **Privacy.** Data is split into chunks, encrypted and stored on different nodes that controlled by various hands. There is no way for node owner to spy data or steal it until data owner disclosure the private key to the public.
- **Low cost.** Utilization of currently unused storage space brings oceans of capacities to market, decreasing price for storage down. Fluence gives an opportunity for hardware owners to monetize their power similar to Bitcoin.
- **Access management.** Fluence provides flexible permissions management model based on proxy re-encryption technology. It allows meeting compliance with [HIPAA](#), GDPR, HITECH, PCI and other regulations.

The cost of storage and transactions will be defined by the market.

Open Data Market

Using Fluence flexible permission management model, data owners can grant permissions to any amounts of their data to other parties. It is implemented as part of the database (low-level API) and gives all power for creating various business models based on data.

On the high-level side, Fluence provides easy to use web-based solution for uploading data, managing datasets, share permissions or sell data. Anyone can publish datasets meta-information, search for available data and trade.

We believe that based on Fluence technology many industry-focused data markets will emerge. Imagine a service that aggregates genomics data from people makes research and sells results to a pharma company. Resulting revenue is shared between researchers, people who provided data and miners where data is securely being stored.

Use Cases

Fluence could have a significant impact on how companies in many industries interact with sensitive customer data. Most demand will come from areas with government regulations about data privacy.

Data Regulation Compliance

Managing of data lakes becomes too wasteful in regulated industries. Controlling access is painful when data is always moving, changing and being accessed by many users. Companies need to securely share data between employers: business analysts, data scientists, and developers concerning current regulations and limitations to data privacy.

Fluence provides flexible data access management that can be compliant with [HIPAA](#), GDPR, HITECH, PCI and other regulations. It can be used as private enterprise secured storage for sensitive data. Fluence access contracts maintain fine grained management between administrative and authorized users.

Data Sharing and Exchange

Currently, many banks monetize customers' data by providing aggregated information to partner's advertising campaigns. Like medical institutions or insurance brokers, they have to anonymize data, aggregate and implement custom integrations with a partner to get value from the deal.

Also, banks, telcos, and retailers are looking for solutions to get more insights from their data, creating personalized and contextualized experiences for customers. They combine data with other vendors using comprehensive integrations because customers' transactions or location data is private and can't be simply disclosed with third-party.

All of this requires a tremendous amount of negotiations, technical integrations, time and resources. Fluence allows reducing frictions using blockchain-powered contracts and granular data access.

Medical Data Exchange

The recent emergence of Big Data in healthcare (electronic patient records, in-clinic monitoring systems, personal wearable devices) and the [open data](#)

[initiatives](#) benefits significantly to both to clinical practice and research. An improvement in patient outcomes, an increased value of services, more efficient ways of working for healthcare practitioners, opportunities for home care using remote and telehealth technologies, and most importantly, personalized medicine, invaluable contributions to scientific research. Better use of data and technology has the power to improve health, transforming the quality and reducing the cost of health and care services. Even though the benefits are undoubtful, it creates significant challenges around storage, security, ownership and secure sharing of medical data.

Sharing data for clinical and research purposes is [extremely chaotic](#). For example, in the UK the initiative intended to enable large National Health Service individual data to researchers and businesses turned into a disaster with massive data leakage (including de-anonymization of data and leaking data to insurance companies) and eventually failed.

At the international level, sharing data is almost impossible to the current situation. For example, government population level epidemiological datasets collected through surveillance systems would be beneficial to epidemiology researches, but as countries remain in control of the datasets collected by their public health surveillance services, sharing even historical population level data remains a challenge.

Fluence has a chance to smooth interactions between medical institutions regarding open data initiatives or simple data exchange. Decentralized database with transparent and secure access management gives unlimited opportunities to save time and cost.

Data Marketplaces

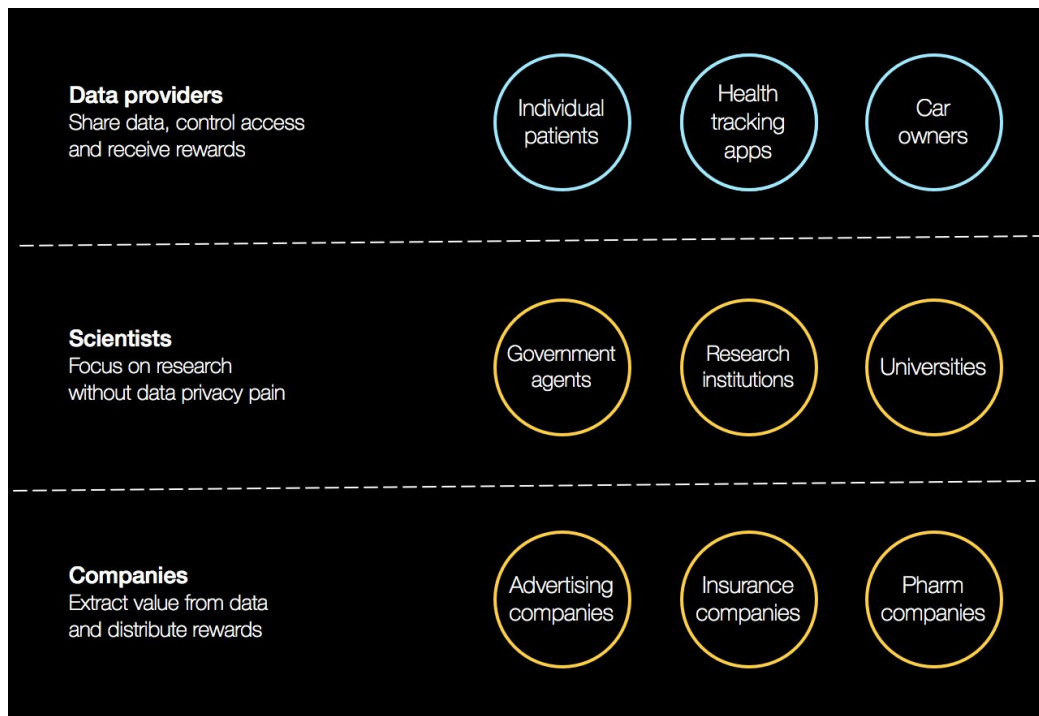
Fluence provides capabilities to create various business models based on customers' data. Apps can rent personal data, create some value and then fairly distribute revenue to data providers.

Nowadays we observe test-drives in many industries that confirm the hypothesis of such approach. There are experiments with automotive, insurance, genomics data.

Toyota Research Institute and MIT Media Lab [partners with BigchainDB](#) to build aggregated data lake for autonomous cars and bring data amounts to the new level. [Encryptgen](#) aims to accumulate genomics information from people to do large-scale science researches and reward people back for that. Such projects require secure and scalable storage for sensitive data.

Fluence provides flexible permissions mechanism that allows anyone to control what is happening with their data. For example, a hospital can decide to do some new research using patients records, but since access control is managed by smart contracts, unapproved data request will be prohibited. And vice versa, patients can grant access to hospital and benefit from researches or data sales.

Another example, in auto insurance, telematics data from devices are being used to monitor driving habits in real time and adjust personal rating. To create more accurate models, companies need to grab private data from a lot of people. Also, telemetry devices on roads report which road particular driver takes and what danger rating it is being graded. Such way of implementing personalized insurance is promising but creates many problems with privacy and transparency. With Fluence, data from many sources could be accumulated in decentralized storage that gives customers full control over how their data was used.



Decentralized Apps

In the decentralized world, apps will use customers data more carefully and customers, in turn, won't share explicit data with apps.

After moving to blockchain based secure authentication systems like [Civic](#), dapps will start to seek databases that guaranteed security. Fluence allows innovative approach for storing data; it has no server that adversary could hack. Because dapp is also hosted inside blockchain without any backend, all user data can be accessed either by the user himself with the personal private key or through smart contract using multi-signature.

Machine Learning and AI

The common way to extract valuable insights from data is running machine learning algorithms (Artificial Intelligence). With increased complexity of processing and analysis algorithms, the demand for computational power increases, forcing to shift from local processors to the cloud. Many areas of healthcare and clinical research would gain from using of cloud computing, but they can't just rent out AWS because of regulations. For example, hospitals use isolated data platforms and electronic medical records.

In future, Fluence will allow running machine learning algorithms on top of encrypted data in the cloud, with full respect for data privacy and security. Researchers will get an opportunity to access massive amounts of data that never had been available before because of regulations and fragmentation. Data owners will be confident that their data has not been disclosed or misused.

Economics

Fluence Token

FLU token is ERC20 token issued on Ethereum blockchain. Its allocation is limited to 100,000,000 and can not be changed. Paying with **FLU** is the only way to rent resources from Fluence network. It is being used as a reward for nodes for providing storage and performing database operations.

To participate in the network, each storage node or arbiter has to deposit some amount of **FLU** tokens on Fluence smart contract as a collateral for honest operations. All operations are processed with providing publicly verifiable proofs (see more in Technical Whitepaper) which ensure that nodes operate honestly.

Business Model

Fluence ecosystem is very similar to decentralized storage projects like Storj, Sia or IPFS. Anybody can create a node, join to network and monetize storage capacities by renting it to data owners. From other side, data owners can put any amounts of data in fault tolerance storage, make query requests to this data and share access.

To prevent malicious behavior of nodes, every operation in the network is incentivized. Typical operations are:

- Add data to storage
- Select some amount of data from storage
- Calculate some formula using data in storage
- Deposit funds to network
- Withdraw funds from network

Each operation is created as the transaction on the internal blockchain, securely verified and stored by nodes responsible for storage of particular client.

Deposit and withdraw operations are performed by special smart contract on Ethereum blockchain. Thus, economics inside Fluence network is entirely in the hands of nodes and clients.

Storage Contract

When client aims to use Fluence services, it creates storage contract and funds it with **FLU** tokens. The minimum funding is 1 GB/month, and default contract length is one month.

Every hour all nodes that store client's database should prove storage and receive a reward proportional to their participation. Every request from a client to data is also billed in accordance to its computational difficulty.

If node violates contract conditions, can't prove storage of database or refuse to perform requests, it loses rewards for all contract period. Otherwise, if all parties are fair, nodes receive contract reward and negotiate contract prolongation.

Sharing Contracts

Data sharing and access management are also implemented via contracts. When a client gives somebody access to data, he specifies a price for reading, writing (if desired) and querying. The client creates a contract with all

parameters and puts it in blockchain along with special re-encryption key, generated by Fluence.

Another party reaches node cluster that store data pay for access as much as specified in a contract, and receive data. Due to an atomic payment mechanism, data buyer can stop paying if data turns out to be fake or counterfeit.

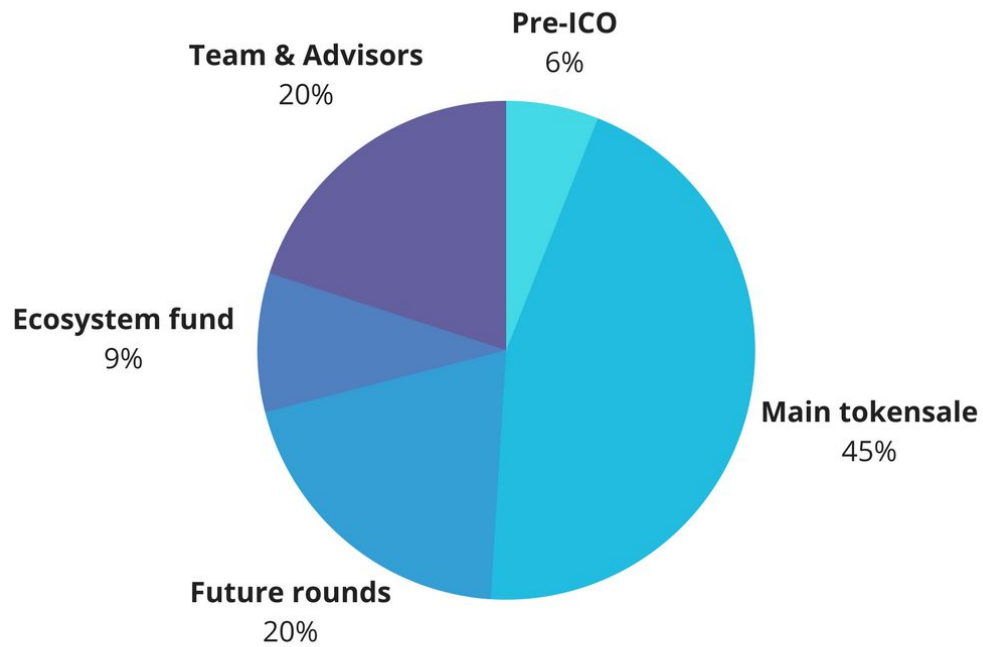
Sharing contract protects all parties that participate. Data owner keeps his data encrypted and gets revenue while buyer makes requests to data. The buyer doesn't pay before reading data. And node receives rewards for performing all operations: storage, writing, and reading.

Tokensale

FLU tokens will be issued on Ethereum blockchain during crowdfunding period. Token creation will take place using Ethereum smart contracts. Investors will send ETH to specified Ethereum address, creating **FLU** tokens for defined exchange rate.

Tokensale participants will be able to send ETH to tokensale address only after the start of token sale period. End of the period will be defined by specified block number on Ethereum blockchain or when targeted **FLU** cap is reached.

FLU token distribution	
Maximum number of FLU tokens	100 000 000 (100%)
Presale (Pre-ICO) issued FLU	6 000 000 (6%)
Tokensale issued FLU	45 000 000 (45%)
Held for future rounds	20 000 000 (20%)
Fluence team tokens	20 000 000 (20%)
Ecosystem fund tokens	9 000 000 (9%)



Budget Allocation

The funding received through the tokensale will be used for development purposes, for building storage ecosystem and partnerships.

Ecosystem Fund

Ecosystem fund is a special multi-signature address on Ethereum network. It keeps funds that can be used by Fluence team for incentivizing network participants and token holders, running developers hackathons and contests, using for other marketing purposes.

Tokens are first time deposited to fund after the tokensale and locked there for 6 months.

Team



[Evgeny Ponomarev](#)

CEO

[LinkedIn](#)

Entrepreneur and product manager, highly passionate about blockchains and machine learning. Built [2GIS Dialer](#) mobile app that became most popular third-party Android dialer app in CIS and processed more than 65 million calls per month. In charge of PM track of [CodeFest](#) developers conference (1500+ attendees/year). Mined his first Bitcoin on CPU [in 2011](#).



[Dmitry Kurinskiy](#)

CTO

[LinkedIn](#), [Github](#)

Software engineer and tech lead. Develops software, builds teams and projects since 2004, including international start-ups. For the last year, he has been [creating](#) a private trading platform, streaming & handling both live (~1k points/sec on a single commodity) and historical (>10Tb/commodity) data. Interested in distributed systems, functional programming, and stream processing.



[Constantine Solovev](#)

Software engineer

[Github](#)

Seasoned software engineer with experience in distributed systems and big data processing. At CleverDATA he developed a highly loaded data exchange platform for RTB advertising.



Alexander Demidko

Technical Advisor

[LinkedIn](#), [Github](#)

Experienced distributed systems engineer with deep understanding of computer science, game theory and recent interest in machine learning. At [Metamarkets](#) he built a petabyte scale analytics platform capable of processing hundreds of billions of events per day.



Michael Egorov

Advisor

[LinkedIn](#), [Github](#)

Michael is a CTO and co-founder of [Nucypher](#), a company providing an encryption layer for popular Big Data frameworks. After obtaining a PhD in physics from Swinburne University he built [ZeroDB](#), an open source end-to-end encrypted database.



Nhan Phan

Advisor

[LinkedIn](#)

Nhan is a VP of Engineering at [Metamarkets](#), where he manages a distributed team of engineers building a highly scalable and reliable analytics platform. Previously he was a VPE at iSocket, which was acquired by Rubicon Project in 2014.

Roadmap

The roadmap is in progress and will be adjusted after further researches.

Proof-of-Concept

Decentralized database

- Store encrypted data
- Query encrypted data
- Connect any node to network

Fluence PoC will allow anybody to store data in the decentralized, encrypted database, query it and manage access. All operations with data will be timestamped to blockchains after nodes consensus. Any node will be able to join or leave the network at any time.

Short term

- Node clusters and arbiters
- Storage contracts
- Sharing contracts

Mid term

- Nodes incentivization for storage and operations

Long term

- Centralized data marketplace
- Machine learning