Decentralized Social Graph Protocol

An open and transparent data interchange protocol for base-layer agnostic and censorship resistant social media.

Document Revision 2020-01-02 Draft

Bevan Barton
bevan.barton@gmail.com
https://www.bevanbarton.com/

Jordan Mack
jordan.mack@brilliantnotion.com
https://www.jordanmack.info/

Abstract

The Decentralized Social Graph Protocol (DSGP) is designed to enable interoperability between different vendor front-ends (VFE) on a singular decentralized graph structure for the purposes of global-scale social media. The protocol places emphasis on remaining agnostic to any base layer protocol or technology. Emphasis is also placed on censorship resistance, but in the form of user choice, not absolute permanence.

The design of this protocol takes into account the lessons learned from the numerous approaches of first-generation decentralized social media, as well as the business concerns of traditional social media. It is expected that VFEs will choose to operate very differently and that this competition will ultimately benefit the ecosystem as a whole.

Motivation

Social media has become a common source of information sharing and a daily activity for billions of individuals around the world. The ability for individuals to have a digital voice is a prerequisite for personal freedom in a world that is quickly becoming more and more reliant on digital technology.

Today's social media is dominated by centralized entities that are inherently subject to undue influence. Individuals and organizations with political or monetary motivations have repeatedly demonstrated that they are able to game the system. This gives privileged individuals the ability to alter narratives and shape the perceived public opinion in a way that is both unjust and damaging.

Decentralized blockchains have proven that they are the most resilient data structure known to man. By leveraging blockchains and other decentralized technologies, we can build a better infrastructure for social collaboration on a global scale.

Document Scope

This document includes opinionated observations of existing solutions and attempts at decentralized social media and will make recommendations for future directions. High-Level protocol design is included, which takes into account these observations.

This document does not address specific strengths and weaknesses in any of the underlying decentralized technology options. It takes a neutral stance without favoritism while simultaneously recognizing that the technology options available are rapidly evolving.

This document does not go into depth about the specifics of the algorithms which would be used for data filtration, or the structures which could be used to incentivize specific behaviors. These are the concerns of the competing VFEs and their accompanying communities.

Goals

- Maximize decentralization and trustlessness while using a common protocol.
- Resist censorship in all forms, while empowering users with individual choice.
- Remain agnostic to all base layers and technologies.
- Remain lightweight and flexible rather than strictly defined and rigid.
- Allow for global scalability both in data creation and data procurement.
- Allow different business models to compete in the ecosystem without favoritism.

Examination

Advantages of Decentralization

Censorship Resistance

Censorship is widespread throughout the world. The effect is not felt equally by everyone, but it is universally experienced. Any platform with visibility has value, and anything with value immediately becomes a target.

Censorship comes in many forms. The most common is by authoritative entities or politically motivated entities. Their deep pockets and control of traditional media give them a large amount of control over public perception. Another common form is through algorithmic ranking. The algorithms are designed to give the user what they want to see, leading echo chambers. Users often inadvertently censor contradictory viewpoints, leading to a distorted perception of reality. A third form is a result of spam and advertising. Any entity with a promotional agenda is able to force its information into view, which comes at the detriment of real content since the time available to sift through and consume information is not unlimited.

Decentralization can aid the situation by providing more resilient and transparent processes than those in use today. By making the total data set accessible by any party, the barrier to entry to become a VFE is much lower. This will lead to more choices available to each user, and a total reduction in malicious censorship.

Proof of Authorship

The asymmetric cryptography methods commonly used within decentralized software are a natural fit for proof of ownership. Data can be signed by the author so it can be cryptographically proven that they are in possession of the signing keys. Combined with a decentralized blockchain, you can now also prove that the signing event occurred at a specific date and time. These work together synergistically to give proof of authorship without reliance on any centralized third-party.

Freedom of Expression

Users are able to express their thoughts on a global platform without boundaries. This allows for greater proliferation of information. Centralized VFEs in some areas may be forced to comply with information censorship policies, but the underlying global system will remain accessible to any party with the will to interact with it.

Freedom to Choose

Users are free to choose between competing VFEs. Each VFE may have a different set of governing rules and algorithms. Some users may even choose to interact with raw feeds without the aid of algorithms.

Some users may opt for a highly sanitized experience. Others may choose a completely uncensored experience. Many will likely opt for an experience somewhere in the middle. The user's ability to choose is very limited today because monopolies exist for each social media niche.

The current centralized services utilize data silos that make interoperability impossible. The ability to globally replicate the complete dataset enables new markets for competing VFEs. In turn, these VFEs enable user choice.

Moderation Scaling

The chore of platform moderation is a continual cat and mouse game. The moderation practices of current generation centralized services have difficulty scaling to address problems quickly or effectively.

Having multiple competing VFEs will enable higher levels of competition for effective moderation. Systems and policies of the VFE will adapt to address the concerns of their user bases. Enabling user choice between competing VFEs will drive effectiveness.

Failed Directions

Direct User Fees

Charging the user to submit content results in reduced participation. Even if the fees are minimal, the costs can easily become significant enough to discourage those with less wealth.

There is a disproportionate amount of wealth within every population and an even greater discrepancy between different areas of the world. Fees inadvertently create a form of censorship favoring particular parties.

The argument has been made that fees encourage better content, but the empirical evidence suggests that the opposite is true. The addition of fees leads to reduced participation, a reduced network effect, and therefore a lower incentive to create quality content.

Fee-Based Spam Prevention

Charging users to submit content has been suggested as a method to eliminate spam. The theory is that attaching a fee to content submission makes it too expensive to submit undesirable content.

However, all available evidence suggests that the opposite is true. The budget of someone with a potential financial gain is inherently higher than someone participating casually.

Over-emphasis on Immutability and Permanence

Making user content immutable and permanent creates an unfamiliar precedent that users are unprepared for. There is often an accompanying thought that anything spoken for the permanent record must be of profound importance. This sometimes leads to unnatural communication which is counterproductive to community development.

Syncing Binary Clients

Clients which require blockchain synchronization provide a barrier to entry, and a UX hurdle that deters the majority of users. There is likely a small segment of users that prefers direct access, but the majority prefer solutions that have a UX comparable with web 2.0.

Successful Directions

Feeless Submission

Costs associated with participation have been shown to be a deterrent to adoption. Content that is signed by the user retains proof of ownership. Anchoring is still required, but anchoring by the user is optional.

Batching together content reduces the total cost per submission. VFEs can batch together the content of multiple users to reduce costs. VFEs may then be able to absorb the cost of anchoring in order to provide a more user-friendly experience.

Custodial Key Management

Management of the user’s private keys remains a user experience challenge. The ideal situation is that the user would manage their own keys through a wallet that allows for Web3 integration, such as MetaMask. However, it has been found that this requirement is enough to dissuade some users from continuing.

Allowing users to operate using keys that are managed for them has been shown to reduce the barrier to entry to levels equal to that of current Web 2.0 online services.

Providing optional custodial key management may be a necessity for wide-spread adoption until key management is a more common and familiar process for users.

Cached Content (Hosted by VFE)

Reading data directly from data storage, be it decentralized or centralized, may not be possible in all front-ends, such as web browsers. Bandwidth may also be a concern on some decentralized node systems if they are unable to provide enough bandwidth on popular pieces of content.

For the reasons above, content that is cached and hosted on a traditional CDN by the VFE may offer a significantly better user experience. As long as the data remains verifiable, there is little downside to this approach.

Untested Directions

Delegated Curation and Filtration

Sifting through content to bring forward what is desired is a never-ending task for VFEs. This task is primarily handled by complex algorithms today. The results of this approach are mixed. The social media giants have been successful in creating avenues for engagement, but these same algorithms are easily manipulated.

Most users exposed to an unfiltered feed would quickly be deterred by the sheer volume of noise presented to them. It is unreasonable to expect users to curate their own feeds.

If users are given the ability to delegate the task of curating to others, they are given a choice that sacrifices control for convenience. Users would subscribe to feeds managed by other users exclusively, rather than curating the feed themselves. Those being subscribed to are free to use whatever means they deem fit to curate the feed they publish for others.

Some curators may opt to publish only their own content. Others will publish collections of content they enjoy. Some will handle the process manually. Others will employ their own custom algorithms. Some will choose to further curate the feeds published by others, leading to the natural creation of hierarchies within the ecosystem.

Both direct and indirect incentives to provide curation can be explored by VFEs in many forms. Incentivization will inherently lead to paid advertising as well as attempts to control information maliciously. This is no different than it is today. The benefit of this system is that the curation process is transparent, decentralized, and accountable. The user ultimately remains in control and has a choice between competing curators. If a curation source fails to deliver, they are easily replaceable by the user.

Shared Blacklists / Whitelists

Similar to Delegated Curation and Filtration, published blacklists and whitelists could be used to combat malicious entities.

With any authoritative system, the possibility of mistakes and exploitation exists. The benefits of this approach are transparency, decentralization, and accountability, all of which are ultimately rooted in user choice. If maintainers fail at their duties or become corrupted, their presence is easily replaceable.

Editing via Appending

Immutability is often confused with the inability to fix mistakes. This is easily solved by allowing editing and deletion through appending (credit: @abcoathup). Content that has already been submitted to the network is immutable, but revisions to existing submissions can be handled through subsequent submissions that reference the original submission. The VFE can then display the updated content in place of the original.

The original content that was submitted cannot be deleted. Some VFEs will likely opt to give users the option to view the original content and all edits available in the form of revisions.

Unresolved Challenges

Data Mining

All the data that exists within a permissionless global protocol is accessible to any party that wants to access it.

Users should be made aware of the associated risks and participate with the knowledge that anything submitted will be made publically available.

It may sound daunting, but this is not too different than the status quo. Data mining is a multi-billion dollar industry. Privacy concerns are rampant. The main differentiator is that data today is housed in massive privately-owned silos. This gives the owners great power and influence, whereas an openly published system removes control of the data from the select few.

Echo Chambers

Social media algorithms on present-day platforms are designed to favor content which will maximize engagement with the user. This is automatically built upon user actions and preferences which are learned over time.

The content presented by the algorithms predominantly consists of subjects and viewpoints that are stimulating to the user. This commonly results in a mixture of content that both reaffirms the

user's current viewpoints, why also providing a steady stream of inflammatory and/or negative content used to keep the user's attention.

This can result in a heavily distorted view of public opinion since only subjects and opinions that drive continued interaction are presented. The user is constantly engaged, not with reality, but with whatever keeps their attention.

Sybil Attacks

One of the largest challenges on current social media platforms is the usage of bot accounts by small controlling parties. It is estimated that bot accounts are so widespread that the numbers may be in the hundreds of millions.

Bot networks of this size can easily distort public perception. These networks are often used to spread misinformation or political agendas.

There are numerous decentralized identity projects being built to help combat this problem. Among them are, BrightID, Idena, Upala.

- BrightID relies on the analysis of graph data (network connections).
- Idena uses synchronized "AI hard" Turing tests (captchas).
- Upala uses randomized in-person meetings.

Cost of Long-Term Persistence

All current and emerging decentralized storage solutions require maintenance fees for long term persistence of data. Social media can generate tremendous amounts of data, and this data has a cost for storage.

Decentralized solutions that rely on permissionless nodes always incur higher costs than centralized alternatives. Elevated storage costs do not mesh well with social media, for which low user costs are expected.

Perspective Monopolies

Transparency and choice are the solutions being presented to the problems that exist with social media, but the possibility exists that attempts to decentralize could exacerbate the existing problems.

If the market decides to heavily favor opinionated VFEs, the result could be heavily biased information dominating visibility. This is similar to what exists in the news industry today. If these biased perspectives are allowed to fly under the guise of "unbiased decentralization", the effect of the distortion could go unnoticed, and the impact could be much greater.

Specification

Basic Structure

A decentralized immutable blockchain will function as the originating point for all graph data. The points where the graph attaches to the blockchain known as Base Anchors. All data in the graph structure will ultimately originate from a base anchor. There can be an unlimited amount of base anchors.

A base anchor will contain a minimal amount of information designed to locate and validate a single data object. A data object may user content, or links to other data objects, forming a tree-like structure. All data objects are cryptographically hashed so that no data in the tree can be changed without invalidating the entire hash tree. This tree is rooted in the blockchain, making all data in the structure immutable.

Blockchains have high storage costs, which is why only minimal data is stored within them. All user content (text, image, video, etc.), is stored in data objects. The data objects can be hosted in any available medium. This can include both centralized and decentralized storage.

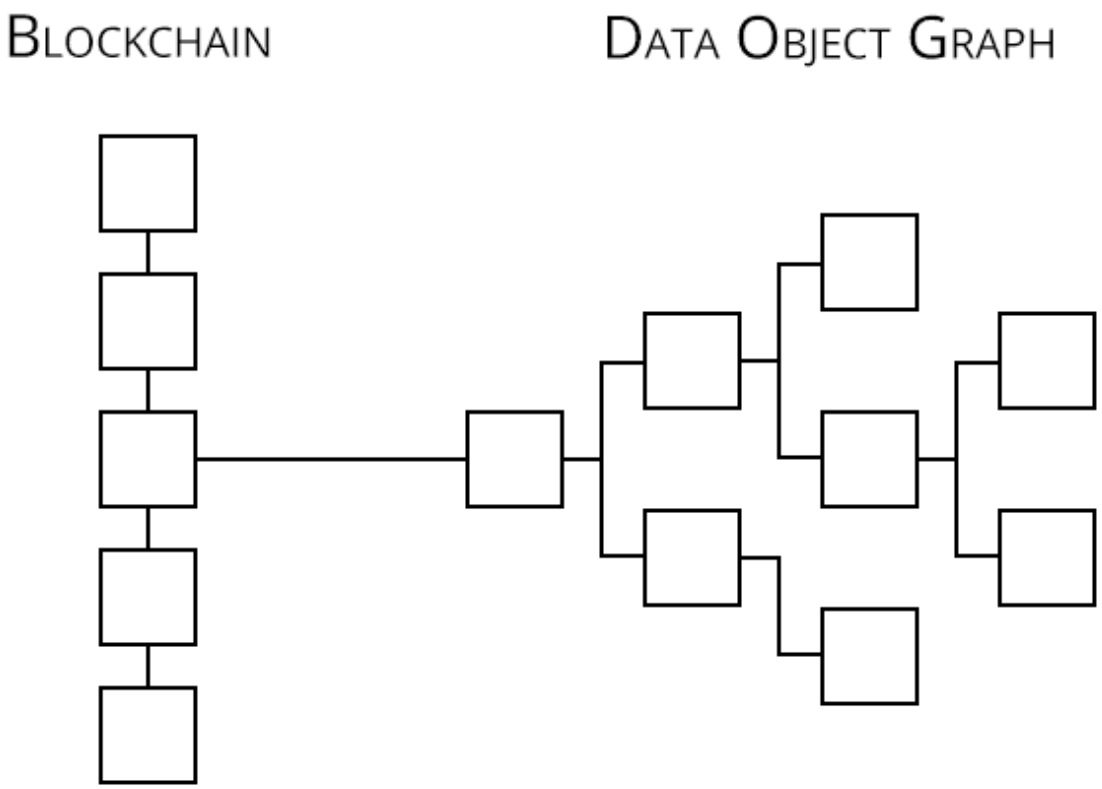

Base Anchors

A base anchor must include the following properties:

- Immutability - Included data cannot be changed at a later time.
- Timestamping - The approximate time of inclusion is provable.
- Arbitrary data persistence - Arbitrary data is allowed to be included and persisted.

Any blockchain that has the above properties is suitable as a base anchor.

Anchor Points

In order to distinguish DSGP data from all the other data, well-known public addresses should be published for any blockchain being used as a base anchor. These addresses, known as anchor points, are determined by the collective consensus of participating VFEs and their communities.

For blockchains with smart contract functionality, this would be the address of a contract that conforms to the protocol. For blockchains without smart contracts, this would be a designated address that transactions are published to, in a format that conforms to the protocol.

If two social media platforms are so different they are incompatible with each other, then they can simply use different anchor points. This gives complete data separation for different types of social media.

Anchor Drop

An anchor drop is the data required to link a root level data object to an anchor point on a base anchor. This single point is the root of a tree of branching data objects containing the data being submitted by the VFE.

The data encoded within an anchor drop must contain:

- Protocol Type: What data storage medium is used?
- Resource Address: Where is the data within the storage medium?
- Media Type: What format is the data encoded with?
- Data Object Hash: What is the cryptographic hash of the data?

Data Objects

A data object is the most basic element in the graph. All points in the graph consist of data objects. All types of data in the graph are contained within a data object.

Examples of types of data that could be included in a data object:

- A user profile.
- A post made by a user.
- Media posted by a user.
- Actions conducted by a user. (Like, upvote, downvote, etc.)
- Relational information with other data objects. (Follows, friends, etc.)

Elements of Data Objects

A data object must include the following elements:

- Identifier
- Timestamp
- Data
- Data Hash
- Owner Address
- Owner Signature

The Identifier should match the hash of the data object itself, whenever possible. In cases where the storage medium does not permit this, the identifier should be a globally unique identifier, such as a UUID.

Elements of a data object that can be reliably inferred by their storage medium should be omitted. For example, IPFS is a content-addressable storage format that uses the hash of the data as the identifier, so the inclusion of a separate identifier should be omitted.

Hierarchical Data Batching

User content must be attached to a base anchor in order to retain provability. This is done by hashing the data and including just the hash in a structured transaction or smart contract.

The costs associated with anchoring a single hash are the same, even if the hash represents a single piece of content, or multiple pieces of content batched together in a Merkle Tree. As long as each piece of content is signed by the user that created it, there is no loss of provability.

VFEs are encouraged to batch multiple pieces of user content into a single tree of data objects before anchoring to a blockchain. This approach has a far greater scale and reduces the costs associated with blockchain storage.

Data Storage Persistence

Decentralized storage solutions are still a work in progress. The performance and cost point for these solutions is still to be determined.

VFEs will likely rely on local centralized caching of data for their users to eliminate performance or synchronization delays that may occur. Data that is persisted externally will primarily be for synchronization with other VFEs, not for viewing by users.

The party that should pay the costs associated with data storage may differ between different VFEs due to differences in their method of operation. However, most VFEs will likely absorb all data storage costs for their users.

Data Permanence

Data may not have absolute permanence. Cost-effective permanent storage and hosting for all of eternity is a fallacy that is sometimes perpetuated by illegitimate projects and the ill-informed. All digital data is ultimately stored and transmitted using the same mediums. These mediums are finite resources with recurring costs that must be paid for.

Anyone with access to the global data feed will have the ability to persist any amount of data at any permanence level they choose. They are free to monetize their actions by creating data markets where the data has value. In cases where data has little to no value, such as mass spam, the data will likely have more limited permanence to a point where it is eventually pruned and lost. The market will decide what data has value and is preserved.

Zero Conf

On-chain data stored in the base anchors should be viewed as the official source of data, but VFEs can opt to monitor and process unconfirmed transactions. This results in data visibility that is closer to real-time.

VFEs may also opt to bypass on-chain solutions in favor of establishing direct connections to each other using proprietary protocols for performance reasons. Data would still be persisted to base anchors, but the direct connections would be used to stream data to each other to reduce lag time.

Network Reliability

The total ecosystem is not limited to any single technology or network. The unlikely event of a widespread consensus failure or loss of data in one technology does not necessarily cause a catastrophic failure of the entire graph.

For additional reliability, all data objects can optionally be persisted to multiple different data storage providers. The root data object can then be bound to multiple different base anchors.

Decentralized systems are inherently more resilient than centralized systems. Combined with multiple persisted data, decentralized systems are more reliable than any single centralized entity.

Scalability

The scalability of the protocol is limited only by the speed at which data can be transmitted and processed. Hierarchical Data Batching makes it possible to persist an anchor drop to any blockchain without impediment. In theory, millions of data objects could be included in a single anchor drop.

Scalability is further benefitted by non-reliance on any single technology or point of failure. Multiple blockchains and data storage technologies can be used concurrently to achieve even higher throughput if required.

Governance Issues

Illegal Content

The handling of illegal content is the responsibility of the VFE to remove from the front-end under their control. This can be accomplished through any of the standard methods of moderation that exist today on centralized systems.

When content is removed from a front-end it does nothing to affect the data that has already been broadcasted or persisted to decentralized mediums. The removal of illegal content on those mediums then becomes the responsibility of anyone hosting it.

The legal repercussions of unknowingly hosting illegal content may vary from region to region. As with all decentralized software, it is the responsibility of the operator to understand the laws that govern their region.

Spam Prevention

The responsibility of spam prevention ultimately falls on the VFE. How this is accomplished is a decision for VFEs to decide. This could include proprietary algorithms, AI/ML, delegation, manual processing, or a mix anywhere in-between. As long as the user has a choice between the different VFEs they use competition will allow the techniques to continually evolve.