

NOTE - Sift protocol Paper:

This document is only used for internal purposes and a work-in-progress. Please do not share, but leave feedback in comments or send them per email (john@mediasifter.co) or get in touch through slack. Thanks!

White Paper.



THE FACT CHECKING LAYER OF THE DECENTRALIZED INTERNET

The Sift Protocol is an open-source framework, which facilitates and incentivised sourcing of evidence for any question of empirical fact. All sourced from a global community of fact checkers and verifiers.

1. Contents

2. Overview

Media Sifter is a news aggregator that excels through superior presentation and content discovery. A key difference to existing, stand-alone news aggregators is that Media Sifter is integrating a second-layer of community content on top of the articles. This can be information to verify, refute or validate parts of the articles, but also relevant content that does not take a side.

This community-driven content creation network is what we call the Sift Protocol and it is a set of smart contracts that aims to source and qualify information from the crowd.

In details, this consists of two steps: 1) source information from the crowd and 2) qualify it.

The first is to fetch information from the underlying participants and have them bring it to the platform. The second, qualify, means to put a "score" on the information sourced to indicate its validity and relevance to users.

Media Sifter requires the Sift Protocol for additional, unique content, to challenge the articles aggregated and as a user-owned discussion platform. However, we are envisioning and designing the protocol to be a stand-alone tool and to become an evidence marketplace. In that case, Media Sifter is the first application and will be used to test and tweak the incentive structures of the protocol, ramp-up its user base and get real world data before it is fully decentralised and operational.

Beyond the first rollout of Media Sifter, our goal for the protocol is to create a community of active contributors that can be directly rewarded for their work. Initially for bringing relevant information to the platform, but later also as publishers on a decentralised content platform. Attracting a healthy reader-base, that understands that content creation should be rewarded, is the backbone of the platform. Once established, we have every intention of feeding back revenues to content creators, including the publications that we will be aggregating from day one.

In the long run, we see the Sift Protocol becoming a fully distributed evidence marketplace, where all types of information can be sourced and traded, both actively (e.g. publishing, selling of information) and passively (pledging, request of specific information). Validity, quality and relevance of the content will be maintained by network participants with aligned incentives and interests.

Media Sifter, the aggregator, is our way of getting valuable content on the platform and giving our users a reason to come and participate from day one. Though we currently design the protocol with a focus on Media Sifter, we keep in mind that this will be an open network that other applications could be benefitting from in the future.

Media Sifter was born in a design environment ([Copenhagen Institute for Interactive Design](#)) and we intend to keep some of the principles for product development. As mentioned in our [roadmap](#), we will be launching simulated versions of the protocol and focusing on the content in the months to come. This will allow for much faster iteration of the incentive scheme and give us more freedom to test different hypotheses. Once we feel more comfortable with the design of the system, we will develop a version and launch it on a testnet, probably some time in Q2 next year.

In terms of on-chain functionality, we are always referring only to the SIFT Protocol. As for now, we do not plan on replicating the aggregator functionality on a distributed ledger. Rather, we are focusing the decentralised parts on the critical aspects of the information sourcing protocol, that center around transparency, trust (or rather not requiring it), incentives and rewards. As scalability and efficiency of the underlying Ethereum-blockchain develops we will also turn up the degree of decentralisation of the Sift Protocol after launch.

Though not visible in written code or a live on-chain application, we are working on the technical feasibility of the protocol and how we will implement it on the decentralised infrastructure stack. A first draft for the architecture can be found at the end of this document, though of course, this is not a technical document, but rather a first version of a concept presentation. At time of publishing it is already somewhat

outdated, as we are solving some of the obvious challenges in the concept and implementing feedback from the outside.

2. SIFT Protocol Process

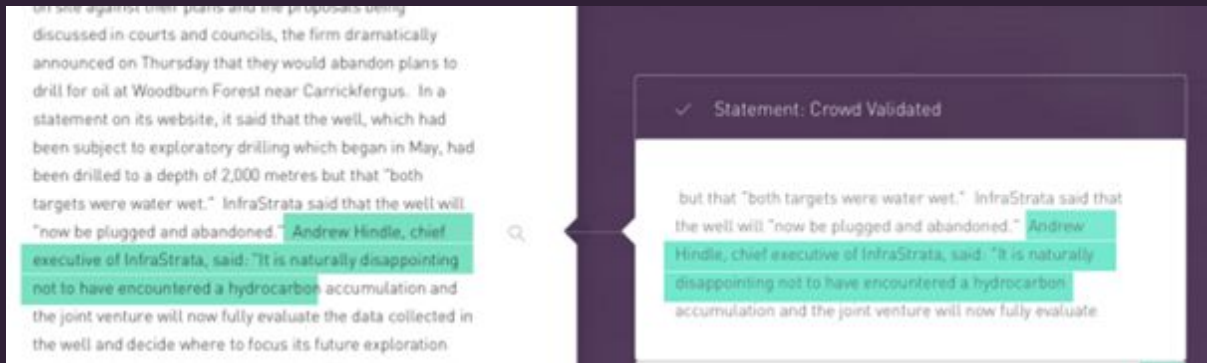
To provide the reader with a better understanding of how the protocol can be implemented, this chapter will walk you through the 3 main parts of the process and use Media Sifter as an example.



2.1. Request for information (Pledge)

The protocol is kicked off by a user identifying a statement that seems so questionable or significantly relevant for the rest of the article that it warrants a reference to establish its validity. The user will then specify the information required in a pre-defined format, categorize it and assign relevance to it by staking SFT Tokens to it (there's a minimum, but not a maximum - this is called the pledge). Once finalised, the request is published on a public ledger and individuals likely to be able to respond are notified. This pledge also becomes highlighted for other readers of the article to show that the community is already investigating the statement and that a bounty already has been posted.

Other users can join by adding SFT to the pledge before the request is resolved. The higher the pledge, the higher the incentive for the community to engage.



Example for highlighting a claim in Media Sifter

The pledge is now open for investigators to provide references to a source that supports or rejects the given statement for 48 hours. Several references can be supported by each pledge, but each reference must go through the review process before the smart contract can be finalised. Once the 48 hour window is closed, no more references can be added, and when the review process for all references has finished, investigators will receive the part of the pledge they are entitled to. The better and the earlier the information delivered, the higher the share of the payout. If the rewards is less than the investigator's stake, the delta will be funded through a payout from the honeypot. While this scenario will not pay investigators additional SFT, everyone involved in a successful contribution will be rewarded with increased reputation (described later in the section on topic-specific reputation) and increase their periodic SFT payout from the honeypot.

If a user believes that additional references are necessary or useful, they may add to- or reopen the pledge by adding additional SFT to the pledge at any time. If the additional pledge is given within the deadline of the first pledge, it will be added on top of that to increase the verification incentive. If the additional pledge is added after the 48 hour window, a new process will be opened with a new 48 hour window.

2.2. Delivery of information (Investigation)

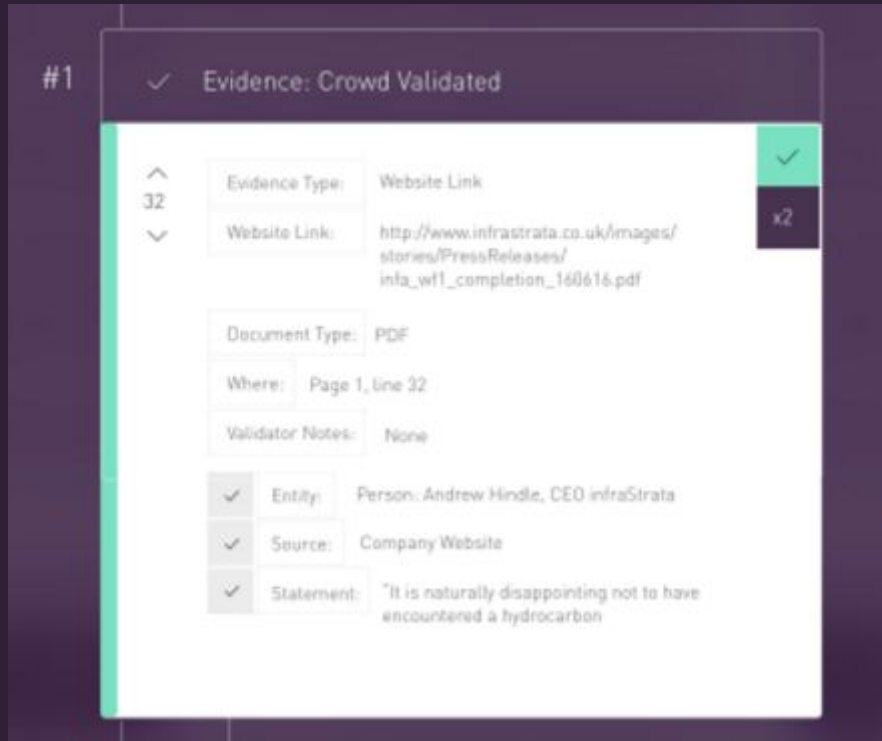
When a pledge has been made by a user, the users with the highest knowledge in the category (as represented by topic-specific reputation) will receive a notification letting them know that they can now add evidence to a statement in the field they engage with most. All users that

see the article will also be able to see that there is an open pledge and will be able to add a reference.

Users that believe they can provide a good reference can take on the role of investigator, source the information required and present it according to guidelines and a pre-defined format. The information delivered can be both from primary (e.g. their own), but also from secondary (e.g. journals) sources. Primary sources cover information such as pictures or video taken of a specific event, official statements from involved persons or companies and in some cases scientific papers that is referred to in the original article. Secondary sources can be news outlets that discuss the same topic or scientific papers that lend credence, but does not specifically validate the statement.

It is vital that the information is presented according to guidelines, to make it as objective as possible for the review process. Relevance and validity are the two decision criteria within the Media Sifter guidelines and it is the investigator's responsibility to convince the reviewer of both. Once presented in the defined format and staked with an amount of SFT that decreases with the investigator's TSR, the information enters the independent review process. This stake is used to discourage spam, and provide a mechanism whereby the investigator can be punished for bad behaviour.

If the information passes the review process successfully, the investigator will be rewarded with the pledge or a share of the pledge if more than one investigator has successfully contributed to the pledge.



Exemplary presentation of evidence within Media Sifter

2.3. Qualification of information (Review)

After an investigator has placed evidence as an answer to a pledge, the community steps up to ensure that the reference actually supports what the investigator claims it does. The structure of the community effort has to be regulated in order to ensure high quality of the results. The Sift Protocol solves this problem through a distributed review process where a group of reviewers are assigned the task of evaluating the reference independently of each other. The review process involves a standard of 7 reviewers, though decreases with increasing TSR of the investigator in questions. Reviewers stake SFT to their vote and the result of the review process is determined by a majority vote of the reviewers. Reviewers that vote with the majority get their stake back as well as an increase in reputation. Reviewers that vote against the majority will lose their stake, which will be distributed pro-rata to the reviewers that voted with consensus.

The review process seems straightforward on the surface but several potential issues need attention to make sure that valid information gets

through, while at the same time maintaining the system secure and preventing spam and malicious attacks. First we will discuss how the reviewers are selected to avoid collusion, monopolies and fraud while fostering speed and quality, and second the exact tasks they are given in the process.

2.3.1 Reviewer Selection

All reviewers have to hold at least a certain amount of SFT to stake their review. This is the minimal criteria. However, it is a user's TSR that will determine their likelihood of being offered a review.

For each reference that needs to go through review, 2 times the amount of required users will be offered a slot as a reviewer. The first to accept and stake the required SFT will get to participate and potentially earn SFT and reputation.

The selection of the potential reviewers is pseudo-random and based on availability, TSR and the size of the original pledge. The goal is to spread the reviewer group as far as possible concerning geography, age, interests (within the broader topic) and potentially race and sex, while lowering the likelihood of collusion. In theory, this should lead to the best utilisation of crowd knowledge while accelerating the reviewer selection.

How the potential reviewers are selected is significant to the final outcome of the review process. However, the number of participants on the network will start out low and is expected to grow over time. To mitigate this, we will introduce developmental stages of reviewer selection in the protocol.

Initially, the Sift Protocol will lack the in-depth information about participants' TSR and the low number of users will make it more difficult to prevent collusion. This should be a temporary challenge for the Sift Protocol, because a small group could potentially dominate the reviews of a certain topic. Therefore, we foresee the reviewer selection to develop in three different stages.

Stage 1

First, users will be drawn randomly from a pool that contains alpha and beta version users. At this stage, we can siphon users into the protocol without giving them the opportunity to become a majority. At this point a lack of TSR is not enough to eliminate any user from the reviewer pool.

Stage 2

As users gain reputation in their respective areas of interest, TSR will become widespread and drive reviewer selection. Users with higher TSR in a category will have a higher chance of being offered a review, but the entire review pool will still contain reviewers without any TSR. Stage 2 will continuously add reviewers based on their TSR as the user base grows and more potential reviewers can be added on this basis. This leads to an ongoing and smooth transition to stage 3.

Stage 3

When the platform has matured and the number of available reviewers with TSR grown, review positions will be offered based on a pseudo random selection weighted by TSR. The more TSR a user has, the higher the likelihood of being offered a review. Users without TSR may still be offered review positions from time to time to ensure new users can participate, but the reviewer roles will be filled by users with high TSR for the most part, to ensure that the consensus reflects what the users that have historically engaged with the topic believe to be the best answer.

2.3.2 The review process

When offered a review, a user will have 60 minutes to accept and conduct the review.

Reviewers will vote on the question of whether the reference answers the pledge or not, and assign a quality score to the evidence based on their TSR. They are also given the opportunity to leave a comment, in order to

support the rating they have given. If more than one reference is added to a pledge, the reviewer score will determine the payout of the pledge and the reputation gain of each investigator. Moreover, the review score will determine the presentation of the reference within the articles in the aggregator.

The review process will result in quality score for every piece of information added to the platform. The quality score helps readers understand the quality of the information and how to interpret these. If the evidence scores above a certain threshold, it will also lead to a payout for the investigator and guarantee a return of his stake. The gain or loss of TSR for the investigator is directly linked to the score.

One might suggest that the payout for reviewers is too small to justify the stake they are asked to put in before they see what they are reviewing, but this assumes that the payout is based on a single event. As the review is based on consensus, users will have a high level of control over how the end result of the process is, and their efforts will in most cases determine if they vote with consensus or not.

2.3.3 Accelerated reviewer replacement

Some reviewers may fail to cast their vote and assign a quality score in time, and will therefore lose their stake and some TSR. Their stake will be added to the reviewer-pool, selected on a similar basis as the standard process. The replacing reviewer will have a guaranteed return if voting with consensus, as the pool will have a minimum of one user-stake.

3. Token & Reputation

As mentioned previously in this paper, we are planning to introduce a token specific to the SIFT Protocol, called the SIFT Token (SFT). SFT has numerous functions on the platform and is intended to be used both as a form of settlement around information as well as access to the core community-features of the protocol.

The second key pillar of the platform is topic-specific reputation (TSR), which is non-tradeable and accessible through contribution and use of SFT. We will introduce both in the following.

3.1 The SIFT Token (SFT)

SFT is a standard ERC20 token and a fixed amount will be distributed early 2018. It is the main instrument of value and is intended to facilitate the transactions on the evidence marketplace that we envision. Its use can be broadly categorized in two groups, transaction and contribution.

The first category, transactions, is a typical use case in the crypto-space, where a platform-specific token can be used to settle transactions. In the envisioned evidence marketplace the possible applications are numerous and we expect other applications to find use cases unique to their services. In the short-term, SFT's are likely to be only used on the Media Sifter news aggregator and transactions include, but are not limited to, the following:

- Request and pay for information, e.g. a successfully completed pledge
- Get rewarded for successful contribution to the platform
- Directly pay contributors and content creators you like (donation-like)
- Participate and earn from review process
- Interact with the community (e.g. commenting, likes)

The second category of use cases are directly linked to a user's contribution, quantified by TSR. We are designing the system with the goal in mind, that reputation (and the income from it) should be of more importance for a large share of users than direct transactions on the marketplace. As we are building Media Sifter (and the SIFT Protocol) to be non-profit, all revenues that will enter the honeypot will be directly redistributed to users, according to their contribution. A higher overall reputation (calculated as a weighted average of a user's TSRs) will lead to an increased payout from the honeypot. In order to gain reputation, a user will necessarily require both SFT as well as contribution that is beneficial to the platform. In the launch-case of Media Sifter, these include but are not limited to:

- Access Media Sifter through monthly subscription¹ paid in SFT
- Get rewarded from honeypot in SFT
- Contribution to gain reputation requires SFT for transactions
- Increase reputation-gain (and hence payout) by staking SFT

Details about TSR can be found in the next chapter, where we discuss how it can be gained and how it will determine the payout from the honeypot. The honeypot itself will be comprised from several sources, which include but are not limited to:

- Monthly subscription fees from users
- Lost stakes from participation, e.g. during delivery or review of information
- (Micro)-fees for comments and likes

3.2 Topic-specific reputation

The second critical part in getting the incentives right is topic-specific reputation. From a technical point of view, this information will likely be stored and calculated off-chain at launch. As with the rest of the protocol, we plan to increase the degree of decentralisation over time, once the underlying infrastructure will allow for increased scalability and lower transaction costs. A mid-term solution might be an approach similar to Colony's recently published whitepaper, where calculations are done off-chain and only results are reported to the blockchain by trusted users. This, as several other aspects of the platform, is yet to be determined and will be gradually improved over time.

As for the application of TSR, every user of the SIFT Protocol will have a score in the categories they participate (e.g. Environmentalism) that attempts to quantify a user's historical contribution to the platform. The higher the reputation of a user within a category, the higher his influence in that particular domain. Reputation is gained through positive contribution to the platform, through e.g. successfully delivered

¹ All but the basic reading-functions of Media Sifter will require a subscription. Active contributors are likely to have payouts larger than the initial fees, effectively only charging readers that pay the content and community for their work

information or voting with the consensus. Reputation can not be bought from the outside and will be non-tradeable to begin with². Is it a number that can grow infinitely depending on a user's contribution and becomes relevant only relative to other user's on the platform (e.g. 200 TSR on Environmentalism is outstanding if the average is 50, less so if it is at 1000). Every user has a "total" reputation score that is the sum of all the child-categories, including a non-topic related score (for e.g. referrals).

The platform will launch with a set number of categories, which we will gradually increase with uptake in users and content-width and depth. These will be structured hierarchically, and reputation gained in a child-category will also be reflected in the parent-category and vice versa (gain in "European Environmentalism" also leads to an increase in "Environmentalism").

To avoid monopolisation of TSR at dawn of a new category, it will only be effective following a period after the launch where users will be able to gain it first. While initially centralised, we intend on developing a mechanism that will allow for user-curated categories with attached reputation, as long as a set of minimum criteria is met (e.g. a number of articles and active contributors).

Overall, reputation is required for 3 parts of the protocol:

- 1) Influence on the platform / meritocracy
- 2) Monthly reward from honeypot
- 3) Tagging (ask the "right" people)

3.2.1 How to get reputation

Reputation is earned and lost by action on the platform. For example, when a user adds a pledge that is picked up by an investigator whose reference goes through the review process, they will be rewarded with a certain amount of reputation within the category. The same applies to investigators who have a reference go through review, and for reviewers

² *Though we are playing with the idea of transferring TSR temporarily to other users, as kind of a stamp of approval from the existing community/vouching for other members. Might be part of a release to be tested at a later point in time*

that vote with consensus in the review process. Most actions on the platform lead to an increase or decrease in reputation and we will publish a more detailed guide closer to launch.

In order to align participants' long-term incentives with the platform, increased reputation can be earned by holding SFT throughout the monthly period. Depending on the average SFT held in that period, users can get a multiplier on their reputation earned from contribution to the platform. This multiplier will fluctuate between 1 for users that have little amounts of SFT, and a number of up to 2 for users that own more significant (based on the current value of the SFT) amounts of SFT³. A user's multiplier will increase gradually from the starting point of the line and be capped at a fixed number of tokens (we will start and test 0.1% of total SFT supply).

The multiplier itself will be a fiscal tool to incentivise SFT holding vs spending. In times of lower demand for SFT, it might make sense to increase the multiplier closer to 2, to incentivise investment in SFT and the platform. In times of low turnaround on the platform, it could be useful to decrease the multiplier to close to 1, to disincentivize holding and to foster SFT use on the platform (a bit like interest rates in the real world).

Exemplary numbers for the multiplier are in the table below:

Action	Pledging	Investigating	Reviewing
TSR Payout	50	100	20
With 1,5 X multiplier	75	150	30

3.2.2 What Reputation can be used for

Reputation unlocks certain user roles. Some pledges will be worth more SFT, either because there are multiple pledgers that want the information, or because a single pledger has decided to pledge considerably more than

³ The extent of the multiplier has to be limited in order to prevent the possibility of users buying power and influence on the platform, so the cap might change once we see the impact on the live platform.

the minimal amount of SFT for a pledge. This means that the payout for successful validation and participation in the review process will be higher. Because users' chance of being offered a reviewer position is influenced by their reputation, users with high reputation will have a higher likelihood of getting in on these reviews. As the platform matures, topic reputation will also be used to determine which users should be allowed to guide other users' news consumption by collecting the best news sources for understanding an unfolding narrative. This is tentatively known as the Editor role.

Another important concept that we are working on is to allow user's above a certain threshold of reputation to act as moderators in their categories and or to bypass the review process, to e.g. allow for live reporting on an event, which otherwise would be impossible with an ex ante review. Malicious behaviour can still be recognized ex post, through e.g. a dispute resolution process that will then require multiple, high-ranking users in the category. Overall, we see the protocol developing from the rather inefficient, randomised review process towards a system closer to proof-of-stake, where a user's reputation on the platform and the right incentives will produce better results with a faster turnaround.⁴

Reputation also has a social value and is used to signal the specific user's engagement with the topic and how much weight should be given to that user's opinion in reviews and how trustworthy and relevant their references are to other users. This means that the contributions from users with higher reputation will be more visible, and hopefully taken more seriously by other users. Their judgement will directly influence the way Media Sifter presents the contributions they are involved with. *Reputation* has a monetary value. Besides the signaling value attributed to reputation on the Media sifter platform, reputation also defines how big of a stake a user has in the honeypot. The more collective reputation from different topics a user has, the higher the payout from the honeypot will be as shown per exemplary figure below⁵:

⁴ Skipping the review process/Decreasing the number of participants will make it more difficult to assign scores to information added, at least ex ante. A potential solution could be to assign a default-score depending on a user's reputation and allow certain users to vote and change the score when the information is already live.

⁵ A logarithmic curve that leads to increasing payouts per user, when moving up in the reputation-brackets.

Honeypot Payout	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%
REP % low to high	30%	20%	15%	11%	6%	5%	4%	3%	2%	1%
Payout example ⁶	3,3	5	6,6	9	16,6	20	25	33,3	50	100

Reputation is not for ever. Each month, after the honeypot has been paid out, the reputation tokens will be cut by 20%⁷. This happens to keep the Reputation up to date with the ever going cycle of news, and ensures that new users have a chance to reach the top level of reputation. If a user leaves the system for a longer period of time they will become gradually less reputable, until they end up where they started (The starting point is still to be determined).

5. Suggested Architecture

This chapter is a first draft of how the architecture of the protocol (and the product overall) could look like. This is by no means a technical description, but rather the initiation of a thought process. As stated earlier, we started working on the theoretical technical implementation of the protocol, but will simulate it off-chain before Q2 2018.

Application Layer: Media Sifter and other platforms

Media Sifter will be the platform that functions as the communicating product of the synergy from all underlying building blocks.

The Ethereum Blockchain will be integrated to the platform through the Ethereum JavaScript API (web3 library) which allows for communication between the two.

⁶ SFT / user, at 1.000 users and 10.000 SFT in the honeypot

⁷ As all variables on the platform, number to be tested, tweaked and adapted.

IPFS will serve the frontend as it works much faster than the normal HTTP protocol. Content such as images, PDF, text etc. will all be stored in IPFS and given a cryptographic hash (more about IPFS in the later section).

The platform will be connected to a preferably decentralized database such as BigchainDB (the go-to decentralized choice for many new blockchain oriented projects) or a traditional relational database. This will contain information related to transactions, pledges, and much more.

Other platforms will be able to implement our SIFT Protocol into their development stack or any of the other solutions which are being mentioned throughout this proposal, as they are all open source (SIFT Protocol will be too as it will be built for on-chain purposes).

Processing Layer: Ethereum Blockchain + Raiden Network + SIFT Protocol

The transaction capacity of the current Ethereum Blockchain is very limited.

Full blocks, high fees and long confirmation times are the known effects of this limitation and the problem is scalability. Ethereum is only able to confirm approximately 10-12 transactions per second. Media Sifter would require 100s of times of this, if we were to use only the EVM for the processing of all logic and the transactions (this is with regards to ALL transactions including the collecting of stakes from multiple participants and the changes made to their balances, for every piece of content added). Additionally, it would be a very slow user experience and expensive with regards to fees (Gas).

Concluding, to operate the most efficient way, an interplay of multiple solutions is required.

We plan to incorporate the Raiden Network into our solution. It uses a mix of meshed payment channels, deposits and cryptographic tricks which will allow *mathematical computations* on the distribution of stakes and rewards to be executed completely off-chain. The Raiden Network will be used to execute procedures such as the collecting of stakes from the pledge process and as mentioned previously, the respective distribution of those. The Ethereum Blockchain will be complemented by the SIFT Protocol and it will allow smart contracts to be used to eventually settle netted changes

which happened off-chain to the balance of users and the Honeypot with the respective hash of the finalized review process. The Ethereum Blockchain will allow users to have a verifiable receipt of changes made to their balance.

SIFT Protocol will work as a complementary protocol to run on the Ethereum Blockchain. The SIFT Protocol itself will be built as an on-chain solution consisting of multiple smart contracts. The smart contracts will consist of unique features which ultimately seeks to spread influential power to users that act in a beneficial way to the platform.

The SIFT Protocol will include *cryptographic primitives* which will be used for two specific purposes. 1) allow participants of the review process to commit their statement in all secrecy while remaining unaware of other participants until the review phase is finalized.

2) Investigators to be discatenated from their provided evidence and their identity remain anonymous till the pledge has been finalized.

The SIFT Protocol will judge outcomes of reviews based on *consensus protocols* and consensus-decision-making logic, which seeks the maximum possible level of agreement or consent from the majority of the participants.

Additional *rules* will be applied to the previously mentioned consensus protocols which will determine the aftermath of changes to the user balances as well as the Honeypot processed by the Ethereum Virtual Machine.

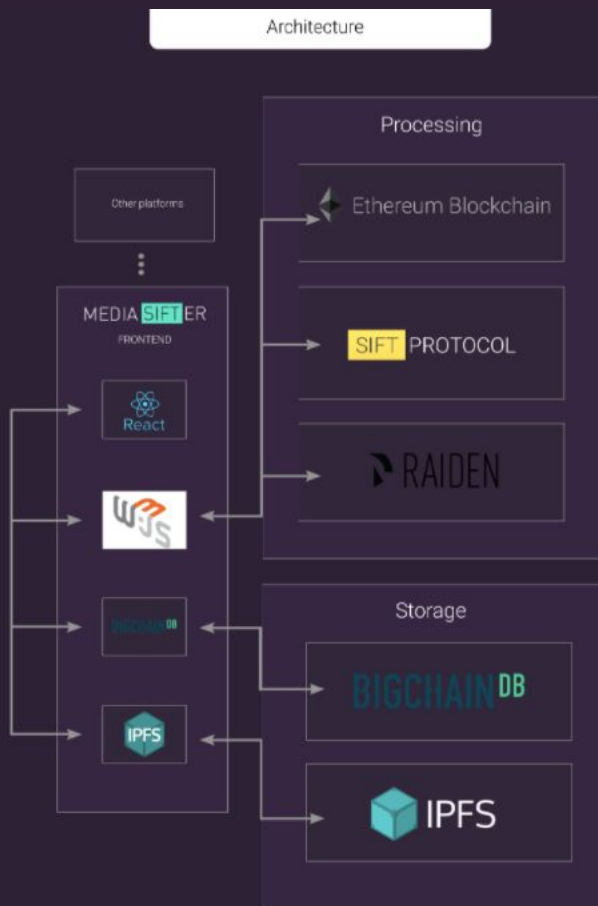
Storage Layer: IPFS + BigChainDB or Traditional Relational Database

IPFS (or InterPlanetary File System) is very useful for blockchain projects, as blockchains are not designed to effectively store vast amounts of data.

Imagine a video smart contract, on Ethereum, for which it would be prohibitively expensive to store the required videos on the blockchain. It would make much more sense to store the videos on IPFS and just store the corresponding IPFS *addresses/hash* on the blockchain.

Instead of referring to objects (pictures, articles, videos) by which server they are stored on, IPFS refers to everything by the hash on the file. The idea is that if in your browser you want to access a particular

page, then IPFS will ask the entire network “does anyone have this file that corresponds to this hash?” and a node on IPFS that does can return the file allowing you to access it.



IPFS uses content addressing at the HTTP layer.

This is the practice of saying instead of creating an identifier that addresses things by location, we’re going to address it by some representation of the content itself. This means that the content is going to determine the address. The mechanism is to take a file, hash it cryptographically so you end up with a very small and secure representation of the file which ensures that someone can not just come up

with another file that has the same hash and use that as the address. Instead of a server, you are talking to a specific object and then you are looking at a path within that object. In the case of Media Sifter, we would use IPFS to store content related to pledges etc. on IPFS, allowing us to allow users to easily verify reasoning for balance changes by attaching a cryptographic hash of the pledge to the transaction.

BigChainDB is a decentralized database and it’s complementary to decentralized processing technologies like Ethereum Virtual Machine, and decentralized file systems like IPFS.

BigchainDB has the architecture to eventually handle the throughput of high-volume payment processors, and directly store contract receipts with other metadata on the database, alongside the actual transaction. This will allow Media Sifter to store any additional metadata to pledges (such as information illustrated on john’s designs) and a lot of other matter.

Some BigchainDB use cases are also like traditional distributed database use cases, except focused where blockchain characteristics can benefit. For example, improving DB reliability by not having a single point of failure, or storage of documents with secure time-stamping.

6. - Appendix

Terminology:

SFT	SIFT Token
CRD	Credibility Token
MS	Media Sifter
Evidence/Reference	the piece of information delivered to validate claims in the articles presented in the aggregator
Investigation	the process where users (= investigators) deliver evidence to validate challenging statement in the articles presented in the aggregator
Investigator	Users who deliver references to verify/unverify/support/unsupport challenging statement in the articles presented in the aggregator
Review	the process where users (= Reviewers) qualify references provided, i.e. assign a "score"
Reviewers	Users who qualify references
Pledging	The process where users (= pledgers) post bounties on information to verify/unverify claims in the article
Pledger	User who posts bounties to get information from the MS-Crowd

Reader	User who consumes MS purely for reading, does not interact with content, community or blockchain (no SFT required)
Honeypot	Central pool of SFT Tokens that is collected in various parts of the platform and paid out according to contribution at the end of the period
Tx	Transaction
TSR	Topic-specific reputation
Editor	Curates topic-specific articles/overviews based on his reputation

5. Open Points

This document is a work-in-progress and far from a bullet-proof concept. As written in the intro, we will test and tweak multiple models for reputation and payouts, SFT stakes, multipliers, timelines, etc in the months to come and will update the published documents with these gradually. A few obvious challenges that we are working on, but have not mentioned above are:

- **Random Voting:** Reviewers stake must be adopted periodically so that randomized voting has a negative average payout. Random voting means voting with the average outcome of reviews, e.g. if in average 70% of all reviews result in a yes and 30% in a no. Therefore the protocol must be designed to make random voting unattractive. Which can be done by making such behavior unprofitable on average. As the platform matures and the quality of the evidence increases, random voting will be increasingly a problem. One of the solutions is to build the system around the reputation (proof-of-stake), rather than randomised review processes.
- **Time-window** to be tested for pledging, not necessarily 48h and constant reset. If window can be reset, an exponential function has to prevent a loop of not-closing the pledge
- **Sybil-Attacks.** We intend on using KYC (e.g. Civic, Uport) and have participants verify identities on the platform. While anonymity is important, particularly in some cases within journalism (e.g. whistleblowing), we believe that trust and the reputation system will benefit from this.
- To have a first layer of content review, we are considering to implement a system that will sort out obviously malicious content and links. This has to be a soft implementation in order to avoid censorship and allow the reviewers and the reputation system to act as the quality gate.
- From a business model perspective, we will run MS for free as long as funds raised and the token reserve will allow, to make the platform as attractive to users as possible. An ideal scenario would be where the token value appreciation and the reserve tokens could fund the platform's operational cost infinitely. A fall back option for this is to charge a

network fee, if required to sustain the platform and the protocol. This is TBD in the future, as the goal is to raise sufficient funds through the token launch to build and run this project without pressure of profitability for the first 24 months.