

Regen Network

System Architecture

Version 0.1
April 4th, 2018

A. Craelius

1. Introduction

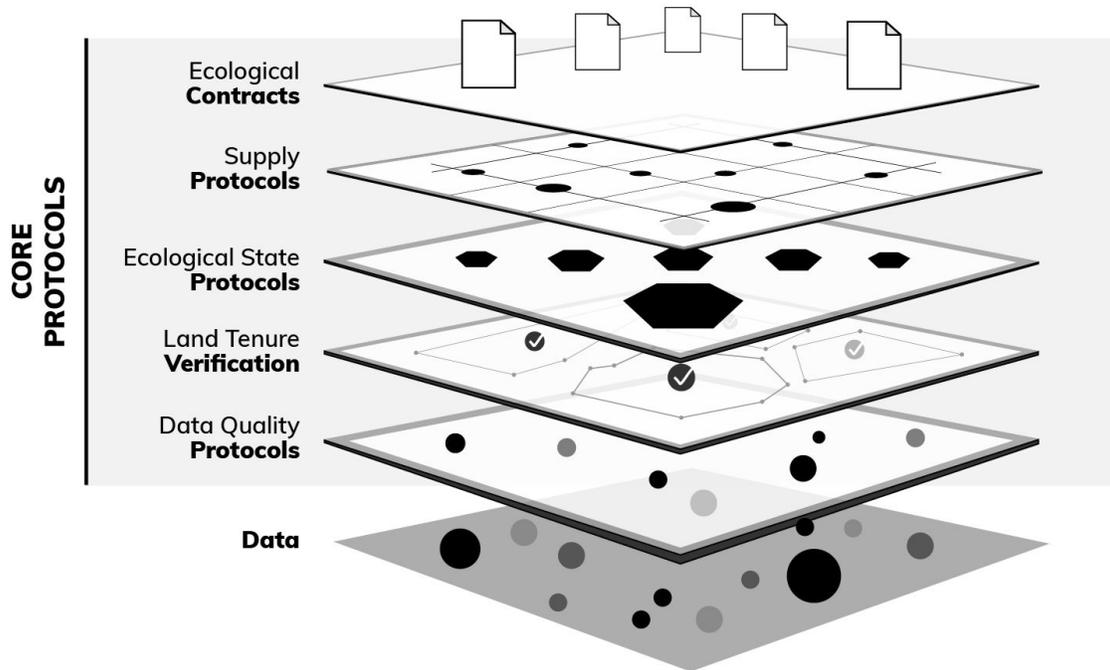
Regen Network's core value proposition is supporting trusted and transparent verification of ecological state and change of state, a streamlined framework for incentivizing this change of state and an integrated system of verified ecological supply from origin to customer. By tracking ecological state, organizations can predict, reward and plan for ecologically regenerative outcomes. This is the core use case for Regen Network technology. Trusted verification of ecological outcomes presents unique challenges. We outline the key architectural decisions of the system and our approaches to solving the most intricate problems. in creating a domain-specific distributed ledger platform focused on ecological accounting.

2. Network Components

2.1. Ledger

Regen Ledger is a domain-specific public permissioned blockchain developed with the philosophy that the most secure way to provide functionality for end users is to code core functionality into the blockchain itself rather than providing a multi-purpose smart contracting language that can be used for anything.

System Architecture



2.1.1. Domain-Specific Approach

With a domain-specific blockchain we can:

- Implement the core functionality appropriate for the ecological domain in a tested and verified manner
- Constrain “smart contracts” to strongly typed domain-specific languages (DSLs) and frameworks which have a limited surface area for bugs and can potentially render contracts visually and in natural human language for clarity
- Ensure high throughput and scalable trustless computation
- Adjust fees to be appropriate for domain users and pre-calculate fees
- Set up governance for implementing domain-specific upgrades as well as emergency hard forks if they are ever needed
- Use a modular protocol-based system to design key elements such as data integrity, reputation, identity, ecological outcomes, etc. and ensure interoperability with other projects

2.1.2. Consensus Mechanism

Tendermint has been chosen as the consensus layer for Regen Ledger because:

- it provides a high-performance Byzantine Fault Tolerant (BFT) consensus layer that allows for any underlying application logic and data layer to be used underneath, and
- allows for a permissioned validator set which reduces overall energy consumption while ensuring high data integrity and federated [governance](#).

2.1.3. Scalable Trustless On-chain Computation

Expensive computations can be scalably performed directly by the compute cluster formed by the Regen Ledger validator nodes. Computations run directly on a distributed ledger are usually of very low complexity for efficiency reasons. In order to perform more computationally-intensive data analysis, such as may be required for Ecological State Protocols, we cannot have every validator node run every computation.

As Regen Ledger is a permissioned ledger run by a set of known validator nodes with a vested stake, we propose the following method to achieve scalable computation as if it were run by every node: the system uses the additional compute power on validator nodes, but only some portion of them run data analysis computations and to guarantee the correctness of these computations, the system uses auditing and challenge windows. Most modern CPUs have at least 4 cores whereas Regen Ledger computations would rarely use more than one core to full capacity. Since validator nodes risk having their stake slashed, and their status as a validator revoked in the case of performing incorrect computations, it is easy to ensure that computations are carried out faithfully by single nodes if our computation is a pure function of known inputs. This follows from the fact that any party in the network could easily challenge the result of a computation and ask other nodes in the network to verify it.

By introducing a sufficiently long window during which a challenge to a computation can be made (and before a computation is valid as an input to a contract) and a long unbonding period to remove a validator node from the network, nodes cannot expect to get away with some malicious computation without serious consequences. Since the types of computations that the platform will perform in this manner will often be analyses of ecological change of state over several months of time, a challenge of a few days is not unreasonable for the intended use cases. In addition to allowing for manual challenges, an automated auditing process will have nodes randomly check each others computations during the challenge window.

2.2. Private Compute Oracle

Regen Foundation will launch and maintain its own private compute oracle. This will enable other organizations to create competing oracles. The purpose of this compute oracle will be to provide a trusted third party computation service for algorithms that must function over private data and where it is not yet feasible to do this in a trustless manner. The foundation will open-source all of the software used to run the private compute oracle, enable other organizations to create competing oracles.

2.3. External Compute Services

The execution of data analysis algorithms is the basis of a significant portion of Regen Network's value proposition and, for the optimal health of the network, computations should be as cost effective for users as possible. To achieve this, the foundation will explore the applicability of compute infrastructures being developed in the distributed ledger space, such as [Golem](#), and look for suitable integration points.

3. Ceres Programming Language

Ceres is the Roman goddess of agriculture and fertility—an apt name for a programming language focused on ecology.

Ceres will be a strongly-typed, pure functional programming language designed to support the domain-specific use cases of Regen Ledger for ecological verification, ecological contracts and supply system transparency. Initially we demurred from creating a new programming language for Regen Ledger, hoping to leverage existing languages where needed. Ultimately, the design goals of our system pointed towards the utility of a domain-targeted language that substantially builds on existing languages.

Ceres has the following design goals:

- Follow a sound, strong typing discipline with optional support for program verification via dependent types, refinement types and/or external provers. In order to create a trusted platform, we aim to ensure trust at the lowest levels.
- Transparently work with stable, globally-namespaced identifiers (such as RDF properties and classes) as first class language features. This will support interfacing directly with the diverse array of data schemas needed to describe ecological data.
- Implement a complexity analyzer that allows for a subset of the language to be provably executable within a certain number of virtual machine operations (as expressed in terms of the function input parameters). This will allow for certain operations to have a pre-calculated maximum cost to execute and thus a maximum transaction fee.
- Allow for a subset of the language to be expressed visually and/or in natural language. This will allow for the financial terms of Ecological Contracts to be clearly stated.

Ceres draws its inspiration primarily from functional programming languages such as [Haskell](#), [Idris](#), and [F*](#), as well as from the domain of logic programming and semantic reasoners. It will be implemented in a strongly typed functional programming language itself (likely Haskell) and will possibly be based substantially on an existing language (such as Idris or F*).

All functions in Ceres will likely be pure (side-effect free) and guaranteed to terminate. Various domain-specific protocols in Regen Ledger will reference strongly typed, pure Ceres functions where the expressiveness of a full programming language is needed. Other parts of these protocols will be expressed as typed data expressions within Ceres, effectively creating mini DSLs for each protocol. The actual execution of protocols, as well as the propagation of any side effects, will be handled by each protocol's execution engine. The goals of this constrained and domain-specific approach are to:

- reduce the surface area for bugs in each protocol,
- simplify protocol creation by facilitating a building block approach of common components,
- and in some cases, in particular for Ecological Contracts, allow for visualization and natural language representation of protocols, as well as a GUI builder.

4. Ecological Protocol Frameworks

Regen Ledger provides three core ecological protocol frameworks:

- Ecological State Protocols (ESPs) define the algorithms and conditions necessary to “verify” a certain state or change of state on a piece of land
- Ecological Contracts (ECs) allow us to fund and reward desired change in ecological state
- Supply Protocols (SPs) allow us to tie ecological state into supply chains in trusted ways

4.1. Ecological State Protocols

An ESP is a specification of the algorithms and criteria needed to verify a certain ecological change of state. A single state protocol specifies a boolean or scalar result on a certain axis. Hypothetical examples include:

- the number of tons of carbon sequestered on a given piece of land in a given time frame into soil and/or above ground biomass
- a score on a scale from 0-10 representing suitability of a piece of land as endangered species habitat
- verification of increase in biodiversity of insects, birds or plantlife
- a boolean true/false value representing whether a piece of land has sufficient groundwater holding capacity to prevent flooding within a given range of rainfall

Ecological state protocols are managed by a curating organization (Regen Foundation, or any third-party entity so inclined), which gives each ESP a unique tag. The curating organization can issue different versions of the same protocol using a semantic versioning identifier. This is very similar to open source software versioning, which gives us a unique tag to identify a given version of an ESP (ex: regen-network/carbon:1.0.3) and also allows us to set version bounds (ex: regen-network/carbon:>=1.1.0). This system allows for protocols to reference each other as dependencies with varying amounts of strictness or flexibility with regards to versioning. It also allows for organizations to gradually upgrade their protocols based on new research and field experience.

The specification of an ESP will be done via a sophisticated smart contracting framework. The long term goal for this framework will be to allow for the complete specification of complex data analysis algorithms in Ceres. This would provide the following advantages over existing data science languages:

- 1) Ceres' strong type system and verification features could ensure that data analysis algorithms are deterministic and wouldn't allow for bugs which could be exploited by hackers or cause the algorithm to fail unexpectedly, and
- 2) Ceres' complexity analyzer could make it possible to pre-calculate the maximum computational cost to run an algorithm for any given set of inputs so that maximum fees can be predetermined.

Since this is a multi-year project that will require significant engineering work to mature, the short term goal of the ESP smart contracting framework will be to provide just enough programming support in Ceres to reference and combine the results of algorithms written in conventional programming languages.

Algorithms in external languages are referenced via a git URL and hash, and consist of source code that will be executed in an isolated container with access only to the relevant input data. At the root of this source code a metadata descriptor file specifies the input data needed for the algorithm to run and the type of the output result (generally a single boolean or scalar value). Since many of the algorithms needed to do ecological verification will involve analysis of satellite imagery and execution of machine learning classifiers, it is likely these algorithms will be programmed in a language like Python or R. It will be up to the algorithm authors to ensure that these algorithms act as pure functions on their input (i.e. return the same result for the same input).

Using the scalable on-chain compute functionality of Regen Ledger, algorithms that reference purely public data can be executed on-chain (or almost on-chain) in a trustless fashion. Some algorithms will need to be run on data that is not available publicly. A basic assumption of our system is that users may not be comfortable making all of their data public. There may be some algorithms that use private data sources such as commercial satellite imagery. To support private data sources, the ESP framework will allow for algorithms to be run on private compute oracles (such as the one launched by Regen Foundation). Each ESP can specify whether or not private compute oracles are trusted and, if so, which ones are trusted. In order to execute an algorithm on a private compute oracle, the data owner will need to authorize access to their data and make a payment for that oracle's services. Functionality to make these payments and to coordinate access will be built into Regen Ledger. In the future, it may be possible to perform some analysis of private data

“on-chain” via technologies such as secure multi-party computation or homomorphic encryption. This will be an area of research for Regen Foundation.

In addition to specifying the underlying data science algorithms, an ESP will often need to filter its input data using Data Quality Protocols (DQP’s) [insert link]. The framework for ESPs will be designed in such a way that DQP’s can be referenced as building blocks.

The basic function of an ESP is simply evaluating state and change of state for a specified area. This can generally be done without actually knowing who is the rightful land owner or steward if reliable, geo-tagged data is available. One compelling application of ESPs is using them as a class of decentralised digital certification (like Organic or Fair Trade), with the goal of promoting good land use practices. In order to link the outcome of an ESP to a land steward’s identity, a Land Tenure Verification Protocol (LTVP) [insert link] will need to be run. Some ESPs may specify what type of LTVP is needed in order for the land steward claim the results of that ESP. A function of the private compute oracle run by the foundation will be to attest that a land steward has attained a given ESP and linked LTVP without revealing the actual land’s geo-coordinates, which allows land stewards to keep their geo-coordinates private if they so choose. In this way, the system provides a feature for preserving user’s privacy while at the same time making it possible to keep ESP results public and linked to the farm’s geo-location so that bioregional-level ecological monitoring still can view these anonymous results. In the future, it may be possible to do this privacy preserving linkage of ESP results to user identity in a trustless way using homomorphic encryption, zero-knowledge proofs, and/or other cryptographic methods.

4.2. Ecological Contracts

Ecological Contracts (EC’s) are one of the core value propositions of Regen Ledger. allow for trusted funding and/or incentivization of specific ecological outcomes. Example use cases include:

- An organization that wants to issue rewards for a specific level of carbon sequestration in a certain region
- A community group that wants to both solicit funds and have them directed to appropriate landowners to support endangered species habitat in a region
- A landowner that wants to request funds to support them to achieve a specific ecological outcome

The EC platform is primarily a smart contract framework for crowdfunding positive ecological change. However, in order to achieve this capability we must also achieve smart

contracting capabilities that make it possible to write EC's for reparations when damage to ecosystems is generated through activity (of course this is not a coercive punishment function, but a system to value ecosystem health whereby two parties would agree that it is in their best interest to value ecosystem health accordingly). In addition to monetary exchange dependent on ecological state there are a wide variety of other smart contract terms that could be used by parties, including ownership, governance and special rights that could be dependent upon a given verifiable change in ecological state. ECs are specified not using a full programming language, but rather via a domain model with lightweight programming constructs where needed. This is to ensure that the meaning of ECs are unambiguous and can be easily presented in a visual and/or natural language form to end users while minimizing surface area for bugs.

An EC is first and foremost constructed using phases. Each phase represents a logical progression in the funding process. The successful completion of each phase is necessary to proceed to later stages, and each stage may or may not involve financial rewards. For example, a restoration project may include an initial phase which simply requires submission of a plain-language description of the specific efforts to be undertaken by the landowner. This initial phase may carry with it a grant to cover costs. A second phase may then involve a reward for reaching a pre-specified benchmark in ecological change-of-state.

ECs will generally reference one or more ESPs and may set thresholds for ESPs with scalar results (ex: > 5 on the biodiversity index). The results of an ESP may also be used to scale reward amounts. For instance an ESP may say that 10 tons of carbon were sequestered on a piece of land and accordingly an EC can specify that \$100 is rewarded for each ton of carbon sequestered and thus a total reward of \$1000. Since smart contracts should be easy to visualize and understand for end users, the "micro-language" for referencing ESP results will include certain constraints to make this possible.

In addition to referencing ESPs, ECs may include other criteria such as submitting plain language text or photos to be reviewed by a specific trusted third party, as well as minimum thresholds for land tenure verification.

ECs define three distinct roles: funders, land stewards, and curators. The curator is always the party that has created the EC, but this party could also be the land steward or a funder. If a land steward creates an EC, then it works somewhat like a traditional crowdfunding system where a land steward is requesting donations. A funder, such as a non-profit, may create an EC to solicit various land stewards to apply for grants from its funds. An entity which is neither a land steward nor a funder could also start an EC to solicit both funders and land stewards as participants. For example, this approach could be taken by poorly funded

local non-profits to make improvements in their community's environment. In cases where a non-profit is the curator, the EC framework will be designed in such a way that funder's donations to the EC can flow through the contract as tax deductible contributions if the non-profit has taken steps to design the EC in a way that aligns with their nonprofit mission.

In part due to the volatility of supply-constrained cryptocurrencies, Regen Ledger intends to whitelist other cryptocurrencies besides the XRN token for use in ECs. These can include stable coins or other coins such as ETH, BTC and the Cosmos ATOM, and will be transferable in and out of Regen Ledger via the [Cosmos Hub](#) and/or [VulcanizeDB](#).

4.3. Supply Protocols

The Supply Protocol (SP) framework builds on top of the Ecological State Protocol framework and adds additional capabilities related to supply systems. One challenge in verifying product origin is verifying that a particular product actually came from the claimed location. A product could be labeled as being organic while actually originating from a non-organic farm. As we are already verifying data on ecological outcomes in our system, an algorithm for estimating expected yield on a given piece of land can be created. It is then possible to create a protocol for verifying supply and to connect that to algorithms for verified ecological outcomes in order to create the possibility of consumer labels based on true, real-time ecological data.

The main additional functionality required for Supply Protocols is the ability to tag products on the blockchain with their origin or en-route geolocations, and possibly photo and weight measurements. These additional pieces of data can then be fed into algorithms that appear quite similar to the ecological state protocols, with the additional function of looking for anomalies in yield. Because an SP protocol may create the opportunity to sell products at a premium price in consumer markets, SP protocol creators may choose to add a transaction fee for the creation of consumer product labels. These could be physical labels that have a barcode which when scanned links to a supply stream and ecological outcome ledger on Regen Ledger.

Because this feature is such a core part of the whole platform, it will launch with a core Regen Network supply protocol that can be tied to any number of ESPs a producer may have satisfied. This core supply protocol will charge a tiny transaction fee for each label that will go directly to fund the network's further development and grants to ESPs. The physical label printed for this core supply protocol may show a Regen Network logo or any logos that are

authorized by ESPs the producer has satisfied. In the case that a supply protocol logo is used on the label, that ESP will receive a portion of the transaction fee to fund its development efforts.

5. Data

An attestation on the blockchain, used to unlock a smart contracted reward for improvements in ecosystem health, is only as good as the data that is used. In order to triangulate and create assurance of accuracy, as well as deter gaming the system, Regen Ledger accepts data from multiple sources linked to the same geographical location, and has several layers of safeguards against bad data (whether it be falsely generated to game the system, or simply data from poorly calibrated or inaccurate sensors.)

In addition to assurances of accuracy and integrity of data, the architecture of Regen Ledger is built to continually incentivize better and more accurate data from multiple sources, and Regen Network as a whole aims to push the envelope on data collection in several key ways.

5.1. Data Sources

This section represents an overview of key data sources that are currently available to compile data and generate knowledge about ecological outcomes for use in Ecological State Protocols.

The data side of the verification protocols consist of four layers:

1. Raw remote sensing data (optical, near infrared, SAR and LiDAR) as well as analyses of that data using vegetative and water indices, both as ancillary data for classification algorithms and to attribute a wide range of characteristics to plants
2. GIS datasets
3. Bioregional sensor networks with large sets of data-points that act as training and validation datasets
4. User-collected ecosystem data (information about soil, practices applied, handheld instruments or other specific data required for the ESP)

More details are provided in Appendix 1.

5.2. Data Schemas

For the Regen Network ecosystem to function coherently, shared data schemas are required. As a basis for algorithms to accept data, they must first know what the data refers to and have it accessible in a comprehensible format. A key development effort in the early stages of the project will be identifying which existing schemas can be reused and which new ones need to be defined. We are following the [W3C](#)'s efforts to create globally namespaced identifiers through efforts such as [RDF](#), and of open data efforts in the agricultural space such as [GODAN](#). Our aim is to support schema development that maximizes interoperability, semantic meaning, and the forward compatibility of identifiers in schemas that evolve over time.

5.3. Data Integrity, Timestamping and Indexing

In order to ensure the trustworthiness of user-collected data (whether it is public or private), users can submit a content descriptor for data stored off-chain that includes the hash of the data, a permanent URL from which to access it, a geographic identifier tying it to a piece of land, and metadata about the data stored at this URL. When this descriptor is submitted to the blockchain it will generate a secure trusted timestamp for the data, which ensures it hasn't been manipulated since its collection date. The metadata descriptor also effectively creates an index of data related to a given piece of land. This then allows verification algorithms to request all data for a piece of land as input—a feature which prevents users from hiding information from verification algorithms. These features enhance the trustworthiness of verification results.

5.4. Data Storage

Most raw data used in the system will be stored outside of Regen Ledger and tracked on the ledger as described above. Essentially, any data storage layer that can be referenced by HTTP could be used (such as [IPFS](#)). Some users may want to keep their data private and hosted on a server they control. As long as the data hosting software implements the protocols necessary to make it accessible when needed for computations, it can be used as a

storage layer. Specific integrations are planned for interfacing with data stored on [Streamr](#) and [FarmOS](#).

5.5. Data Quality Protocols

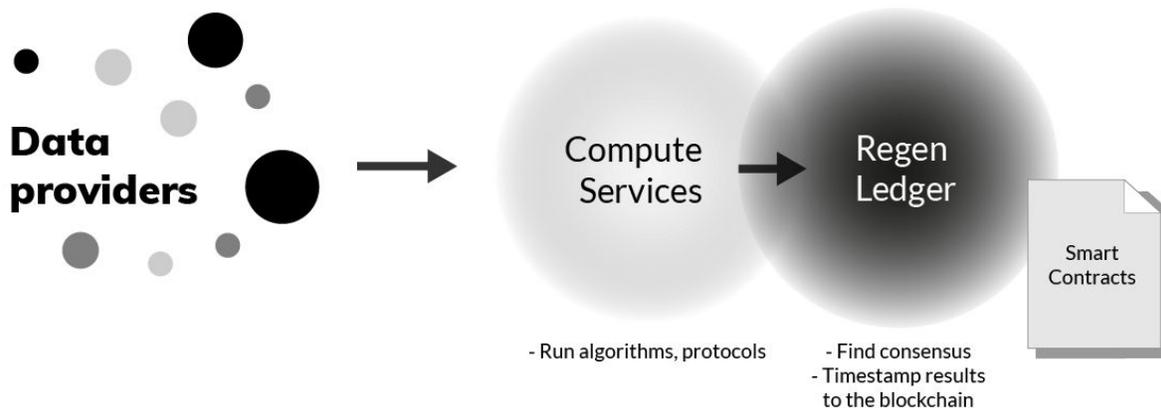
The Data Quality Protocol (DQP) framework allows for a structured way of assessing the quality of input data and can be used by the ESP framework as a way to filter input and/or qualify it with a confidence score.

One of the most basic Data Quality Protocols (DQPs) will be to check that all data has been securely time stamped on Regen Ledger or possibly other blockchains using a standard such as [Chainpoint](#). This is to ensure the input data is tamper-proof.

More sophisticated Data Quality Protocols (DQPs) will use more implicitly trusted data to ensure the consistency of less implicitly trustworthy data. For example, a drone image quality protocol may be implemented, that takes public satellite data and national weather information as implicitly trusted data sources and then runs image analysis algorithms over user submitted drone imagery to look for anomalies. A similar algorithm could be used for different types of IOT sensor input. These Data Quality Protocols (DQPs) would then allow ESP algorithms to assume that user-submitted IOT sensor data is reasonably trustworthy as long as it doesn't set off any alarms in the DQP anomaly detectors. With this model, we have a mechanism for using user submitted data as a finer-grained enhancement of public trusted data. It is important to note that for these anomaly detectors to work well, blockchain timestamps are necessary as a prerequisite. Without these, it would be easy to use AI to generate reasonable looking fake drone footage or sensor input based on available satellite data. In order to trust drone footage and sensor input, the data will need to be time-stamped on the blockchain before public trusted data could have reasonably been acquired.

An additional type of DQP that could be created is a user rating system of data quality. This would allow for more fuzzy, human-based assessment of data streams and could include ratings on multiple scales such as consistency and usefulness. Whether or not this system is useful for a given ESP is a larger discussion, but the intention of the DQP framework would be to allow for even this type of data quality score using the same API.

Data flow



5.6. Data Marketplace

One ancillary function of Regen Ledger will be to coordinate conditional access to network members' private data by other network members possibly via a fee system.

Data Access Protocols (DAPs) allow users to create a contract specifying the conditions under which they will grant access to their data.

Recall that a preliminary step to requesting verification of ecological state via ESPs is timestamping user collected data onto the blockchain. This effectively creates an index of ecological data in the network. As discussed, users may want to keep some of this data private on their own servers and then conditionally grant access to private verification oracles, or as mentioned here to other users on the network. Through DAP's specifying fee structures, Regen Ledger can effectively be used as a data marketplace coordinating access to data. When a user provides the necessary payment for another user's data, a transaction will show up on the ledger reflecting this payment and the public key credentials of the now authorized user.

Servers hosting this private data can then refer to the payment records on the ledger to determine if the user sending a signed request to their server should indeed be authorized. Users purchasing data can have guarantees regarding the integrity of the data they are receiving via the timestamped hashes already on the ledger and possibly regarding its quality via Data Quality Protocols. The DAP framework is intended to be sophisticated enough to support both one time payment and subscription access models.

Regen Network aims to make it easy for farmer data to be co-listed on other existing blockchain data marketplaces (such as Streamr), as the more options for revenue farmers have, the more likely they will be to uphold high quality data. This streamlined co-listing process could be handled via Eco-Apps or possibly directly through DAP protocols.

6. Supporting Ledger Functionality

6.1. Identity

One basic feature needed by many components of the system is making attestations about a user's identity. Regen Network will interface with platforms such as [Sovrin](#), which are dedicated to identity, by designing Regen Ledger to use the [DID specification](#). If needed (for instance, to better support land tenure verification services), certain basic identity support features will be built directly into Regen Ledger and its ecosystem.

6.2. Land Tenure Verification

One of the most difficult challenges in certifying or rewarding ecological change of state is verifying that the party claiming a reward actually has land tenure rights to the piece of land in question. The real-world difficulties with proving land rights carry over to Regen Network. Protocol curators will need to take the necessary precautions to ensure that they have

sufficiently screened participants. The level of verification needed will likely vary depending upon the rewards at stake. Verification procedures will vary by locale. In most cases, the work of doing land tenure verification will need to be off-loaded to a third party organization and ESPs and ECs will need to specify the level of verification required and constraints on which third party verifiers can be used.

To support a diverse array of verification providers, Regen Network will specify a standard API, called a Land Tenure Verification Protocol (LTVP), for verification providers to implement in order for them to be referenced in ESPs and ECs. The foundation will also develop relationships with one or more third party verification providers and steward their implementation of this API to bootstrap the ecosystem. These third party providers may be existing title and or KYC (Know Your Client) companies. A common verification procedure may end up being similar to a standard KYC process with the requirement that additional land ownership/tenure documents be provided and that the verification company maintain copies of them for a certain window of time, as this would allow any disputes to be handled as legal matters in the appropriate jurisdiction. Other verification procedures may involve proof of blockchain-based title in places where this gets implemented or even a proof of location protocol ([insert links to FOAM, etc]).

6.2.1. Organization Management

The basic design of ESPs, ECs, and SP's stipulates that organizations will be responsible for issuing, maintaining, and versioning protocols. In order for this to function, Regen Ledger will have the core concept of an organization, and organizations will be able to specify their own decision making rules. Generally the on-chain decision making support for organizations will involve specifying the threshold of members needed to approve various decisions such as issuing a new ESP version or spending funds in the organization's wallet. This effectively gives organizations built-in multi-signature wallet support.

6.2.2. Token Issuance

It will be possible for organizations to issue their own tokens on top of Regen Ledger that allow their supply to be managed by one or more ECs. This allows for the creation of digital asset tokens directly backed by Ecological State Protocols. The minting mechanism for such tokens is called Reverse Mining, because unlike tokens which are "mined" through an energy intensive process such as Proof of Work, these tokens would represent a net positive ecological impact. Tokenizing living capital assets can be a way to bring value to whole watersheds or bioregions, or to tie value to the health of soil. The tokenization trend which stretches from physical assets, to company shares, and to human attention, will now be able to extend to biological living capital as well. The idea is to create assets that allow people to

buy and sell futures that speculate on the regenerative potential and future health of a whole living system.

6.2.3. Key Management

One of the issues that is most challenging for new adopters of digital assets is the secure management of keys. We acknowledge that this is a complex issue and that it requires thoughtful and careful solutions. For Regen Network to effectively function and be accessible to populations which are not tech savvy (which will likely be a significant portion of our target user base), it is essential that Regen Ledger makes key management both easy and secure for these users. We intend to leverage other projects in the ecosystem trying to address this, such as [Sovrin's DKMS](#), and to create our own solutions if necessary.

6.2.4. Arbitration

There are several places in our system where some arbitration mechanism may be necessary. There could be some to resolve disputes regarding an Ecological Contracts or poor service from a data provider contracted through a Data Access Protocol. We intend to create mechanisms within all of the protocol frameworks to specify an arbitration mechanism (such as [Aragon](#)) that can be applied used in the case of a dispute. In the case that an arbitration mechanism is specified, entering into an on-chain contract would involve agreeing that any disputes are to be resolved by the specified arbitration provider.

7. Eco-Apps

A number of frontend apps will be needed for the Regen Network ecosystem to flourish. Some of these may be existing applications such as FarmOS (which would provide data for ESPs), while others will be created from scratch to support the development, management, and fulfillment of ESPs, ECs, and SPs. These apps are referred to as Ecological Apps or “Eco-Apps” for short. It will be a primary Regen Foundation function to support the development of these apps.

8. Bibliography

Cosmos. “Cosmos/Constitution.” *GitHub*, 2017, github.com/cosmos/constitution/wiki