

TECHNICAL REPORT OF TOKENIZATION PILOT

Overview

SCOPE

First stage of the Elering pilot

Conceptual approach for the test

CONCERNS REGARDING CONSUMPTION DATA

CREATING SMART CONTRACTS

Analysing the dataset

DATA PREPARATION STEPS

Analyzing data points

Analyzing gas cost

Optimizing the dataset

COMPRESSION

Wallets and tokens

CREATION OF TOKENS AND CALCULATING TOKEN SUPPLY

The Test

Takeaways

Conclusion

Next Steps

OVERALL VISION

FEATURE OPTIMIZATION

MARKET DYNAMICS

SECURITY

TECHNOLOGY CHOICE





Overview

Major technological improvements in renewable energy technology and reduced construction costs over the past few years have led to a rapid development of solar and wind generation all across the globe. Unlike the big coal and gas generators, renewables were much faster and cheaper to build, sparking the rise of many privately owned generators joining the energy market and forming the foundation for a new distributed energy system - a system where energy buyers would not need to rely on one single source of electricity but could potentially pick and choose the producers that would fit their needs the best and transact with them directly.

The first sign of this new emerging disintermediated system has been the rise of corporate Power Purchase Agreements (PPAs). Popularised by blue-chip companies in the United States, these arrangements were structured around a transaction whereby a large corporation would commit to buy electricity directly from a new generator over a long period of time (typically 10+ years) for a fixed price. This helped generators to gain financial certainty and businesses to lock in low electricity prices offered by new renewable technologies.

Blockchain as a distributed and immutable ledger seemed like the perfect fit for this new emerging distributed system which is why WePower has designed a platform using blockchain to enable direct transactions between electricity users and producers. We have, however, taken into consideration the scalability limitations of most blockchains - ensuring these factors will not compromise the quality of our technological solutions nor the scale of our business.

Transactions in electricity markets are settled in hourly, 30 minutes or, in some markets soon, on a 5 min basis. This means that to process transactions between a single consumer and producer would create at least 8,760 data entries for a year and if you take a 10-year tenure of a standard PPA this would grow to 87,600 data entries at the point when the contracts are issued. In order for this system to reach meaningful scale, it needs to support thousands of buyers and hundreds of generators.

To build a scalable and cost-efficient technology stack we have set up an R&D team with the main focus on testing and optimizing our technology on different blockchain platforms. This paper describes our first test on the Ethereum blockchain. Tokenization process involves converting a realworld asset, in our case energy, into a digital token that can be stored, monitored and traded on a blockchain. The test used Estonia's nationwide electricity consumption and production data obtained from Estfeed, the central electricity data repository in Estonia managed by our co-development partner Elering.

This test has been conducted using both Ethereum main and test networks. It was carried out across from modelling to live processing. It covered 67,017,042 data points. The following parts cover the technological context, hypothesis, test and implications of this first of a kind blockchain experiment.

SCOPE

WePower has several layers in its technology stack that have to work in sync with our business mission and development roadmap. With classical, fully centralized solutions the estimations regarding system operational costs and performance can be easily aligned with expectations or at least use known practices to assess these. For hybrid solutions that integrate a centralized platform with a decentralized public blockchain, operational cost and performance reference data is not readily available, therefore further analysis needs to be conducted.

For the purpose of the tokenization test, the WePower solution can be described in the following layers:

- WePower cloud platform
- Asynchronous Integration Layer (AIL)
- Ethereum Main network



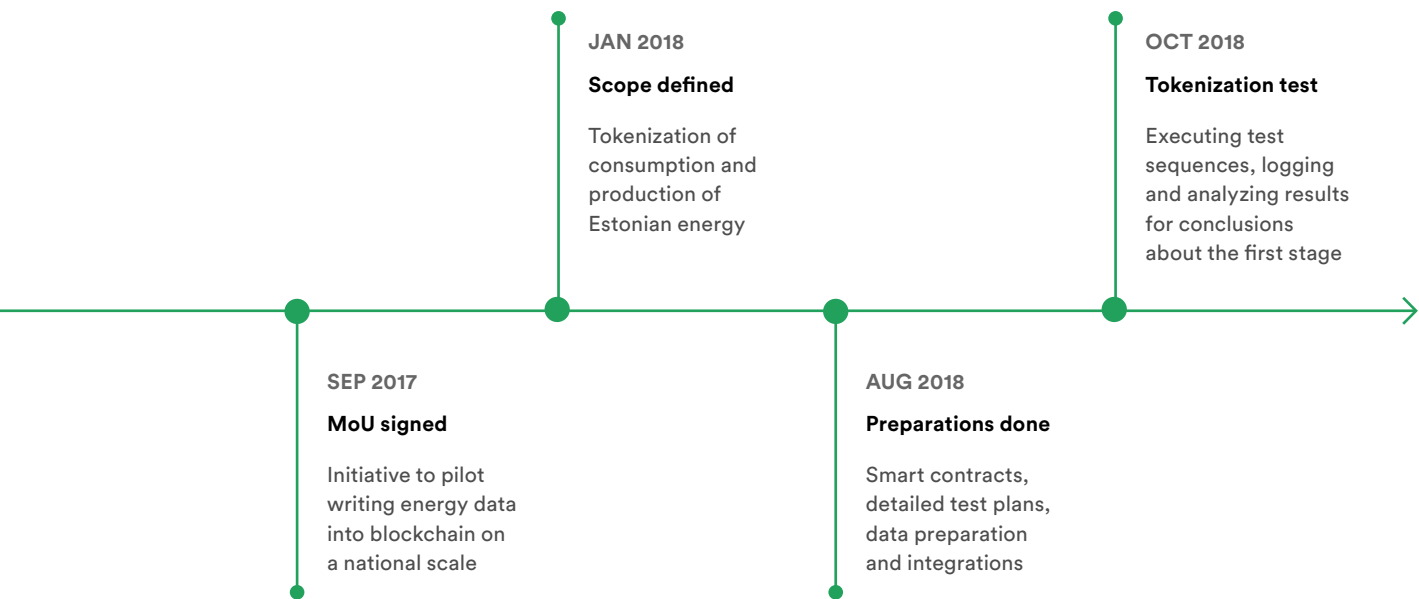
Optimizations for both performance and cost can be achieved at each of these layers. The platform will be sending and receiving data to/from AIL. From there, it is the task of AIL to interface and communicate with Ethereum Mainnet, providing flexibility in the future to consider alternative blockchain solutions or multi-chain architectures. This is where the blockchain/DLT related system operational costs could be optimized, such as the amount of gas used for data storage. There will be two-way communications in the fully implemented end-to-end solution.

In the first phase of the tokenization test, the writing of the energy consumption and production data to the Ethereum blockchain was tested. This is the data set that WePower will utilize in its business activities and where a feasibility estimate is needed.

We analyzed in detail what to write on the chain - as well as considered the cost and time regarding the blockchain activities. This resulted in two layers of activities for optimization purposes, where we optimized the data off-chain and wrote the results to the chain. In the following chapters, we will focus on the write operations and the resulting findings.

The first stage of the Elering pilot

Figure 1: Elering pilot timeline for the first stage



The Elering pilot conceptually started in September 2017, when the Memorandum of Understanding was signed formalizing the initiative to pilot the uploading of national scale energy data into a public blockchain. By January 2018 it was clear that the scope should be the tokenization of electricity consumption and production data. To reach that goal, Smart Energy Contracts were written, detailed concepts and test plans were composed, data prepared and necessary integrations finalized. In October 2018 the test itself was finally concluded. All information about the process is presented below.



The conceptual approach for the test

A Smart Energy Contract (SEC) is a technical term that refers to the entity that defines and facilitates the Smart Energy Token implementation process. A Smart Energy Token is the part of a SEC that represents the ownership of rights and obligations of the energy token holder under a particular power purchase agreement. A wallet is an entity that aggregates ownership of Smart Energy Tokens in one or many SECs.

- SECs contain energy tokens
- Each token refers to a specific electricity consumers' wallet
- Wallets contain references to SECs
- The platform has a capability to issue commands to SECs based on a business logic.

ERC20 and ERC721 token standards have their limitations for enabling the specific transactions needed by WePower. The reasoning for deviating from the well known token standards will be elaborated further in future analysis documents where we dig deeper into the wallet structures, contract versioning, and log chaining along with Smart Energy Tokens and Smart Energy Contracts. We are going to use this concept for tokenizing virtual power purchase agreements, which will become smart derivative contracts¹.

CONCERNS REGARDING CONSUMPTION DATA

The concept of the first stage of the test was prepared on the basis of having the production and consumption data as granular as possible without the risk of not complying to The EU General Data Protection Regulation (GDPR). The GDPR applies to “personal data”, meaning any information relating to an identifiable person who can be directly or indirectly identified in particular by reference to an identifier. The goal of the first stage of the pilot was to generate Smart Energy Tokens to represent energy ownership and match consumption with production. The data used was an entire set of Estonian electricity production and consumption data received from Estonian Transmission System Operator Elering for the year 2017.

¹ International Swaps and Derivatives Association (ISDA) has been working on the Smart Derivatives Contracts concept and issued several white papers on practical and legal matters to use smart contract technologies for tokenizing derivative contracts; Whitepaper Smart Derivatives Contracts: From Concept to Construction.

CREATING SMART ENERGY CONTRACTS

In detailed concept preparation, we had a number of assumptions that had to be addressed. Storing the data on blockchain for transparently representing source and ownership was a necessity. To simulate different energy sources we had to generate three different Smart Energy Contracts.

The first one was a contract with all the consumption and production information. The second one was for fixing the amount of green energy used by consumers. This is recognized in the context of the current test as a “green energy token”.

The third one was for fixing how much was consumed from dirty sources. For simplicity and due to lack of easily accessible granular data, this was decided to be recognized as an “unknown energy token”. “Unknown” because it is not monitored on an hourly basis in a publicly available format of how much and what source of energy is used.



Analyzing the dataset

Our analysis is based on real-world electricity production and consumption data. We worked with both renewable and non-renewable energy production data. In order to retrieve the data from Elering, we utilized Elering’s Estfeed data exchange platform. The Estfeed platform facilitates secure and authorized access to energy (e.g. electricity and gas) meter consumption data by enabling end consumers, scattered producers and network operators to choose who to share their data with via the online platform e-Elering. In addition, public energy data can also be received, e.g. electricity prices or weather information through the platform.

The data provided by Elering was anonymized by aggregating together the hourly consumption data from all meter points within each postcode. Essentially, the data of all neighbors was added together, so that no single neighbor could be identified.

While aggregated data can be obtained for the whole country for a service fee, personal, meter-point level data, can only be obtained for those houses and enterprises, whose owners have agreed to share data with you. This feature, however, was not used in the current test.

DATA PREPARATION STEPS

Analyzing data points

The initial plan was to create the dataset based on one year's electricity consumption of 700,000 households in Estonia. This would have meant roughly:

$$700,000 \times 365 \text{ days} \times 24 \text{ readings a day} = 6,132,000,000 \text{ data points}$$

To calculate the time needed for the test period, we had to consider that the average block time on the Ethereum Main network is 15 seconds. That was the absolute best-case scenario and we additionally assumed we could store 200 data points per block to the contract. In the real world, other people would be able to utilize the blocks or WePower would have to pay the higher gas price to get the full capacity of each block.

The calculation for the time needed was based on the following formula:

$$\frac{\text{Count of data points}}{\text{Data points per block}} \times \text{Seconds per block}$$

With the full set of granularly separated data points, it would have meant that simply storing the data would have taken roughly 14 years.

$$\frac{6,132,000,000 \text{ data points}}{200 \text{ data points per block}} \times 15 \text{ seconds} = 5,300 \text{ days}$$

Analyzing gas cost

Additionally, the gas cost had to be calculated considering 6.132 billion data points with idealistic 200 data points per block, estimating to $6,132,000,000 / 200 = 30.7$ million transactions. Considering the gas consumption per transaction, that would have required $30,700,000 \times 21,000 = 644,700,000,000$ gas, estimating to 5,472,810,000,000 Gwei at 8.5 Gwei gas price level, which can be calculated at any time with the current prices.

For storing of those data points $6,132,000,000 \times 20,000 = 122,640,000,000,000$ gas was needed, estimating to 1,042,440,000,000,000 Gwei at 8.5 Gwei price level. At the time of our calculations, this would have amounted to a cost of 1,042,440 ETH + 5,472 ETH, which was at the time of calculations respectively EUR 208.5 million + EUR 1.1 million = EUR 209.6 million.

Obviously that kind of time frame and cost level is not something we would want to run our test on. Therefore alternative approaches and additional optimizations had to be considered. We needed to compress the data or optimize the dataset so that the amount of data points would be reduced.



Optimizing the dataset

For further optimization of data, we had multiple options. In the end, we decided to compress and reduce the data volume that was to be uploaded for various reasons.

For loading the consumption and production data into the ETH Main network we chose to aggregate the data on a zip code level instead of a granular consumption point level and summarize hourly consumption into monthly consumption. This led us to the volume of $3,837 \cdot 12 = 46,044$ data points. The decision to aggregate the consumption data was also supported by concerns for personal data confidentiality from the GDPR perspective.

As the cost was not the issue when loading the contracts into ETH test network, the data set was prepared on the zip code and hourly consumption and production level, which meant the data set of $3,837 \cdot 8,727 = 33,485,499$ where 8,727 is less than 24×365 due to removed data anomalies. The main constraint for going more granular was the time needed for loading the contracts into the ETH test network.

As mentioned above, basic checks for data quality were also conducted. This led to removing one day's consumption data because there was missing data for some hours, which would affect statistics aggregated on a daily level. There was also an occurrence of abnormal and unrealistic consumption during one hour, which was removed from the data set. But in conclusion, the overall quality of data was satisfactory.

For creating green energy tokens proportionally we considered the total amount of electricity production compared to wind energy production for uploading to ETH test network.

COMPRESSION

Data loading to blockchain was additionally optimized by compressing the data.

EVM (Ethereum Virtual Machine) uses 32 bytes per word which means numbers in EVM and Solidity are internally represented by 256-bit integers. The data we used (consumption and production data with added information) however needed a lot less precision. As described in the section above, for a large scale dataset we had a high amount of data points, but inserting our data into a 256-bit number in Solidity, we were able to accommodate several units of energy data into 1 number and that resulted in less separate data points for the ETH blockchain. In our case, we were able to fit 15 units of data into 256 bits. Therefore we were also able to upload the same data with 15 times less gas cost, 15 times fewer transactions and time.



Wallets and tokens

There were 3,839 wallets generated for the test. 3,837 wallets for customers, 1 wallet for WePower and 1 wallet for the generator. Customer wallets represent ownership of energy tokens and are needed for enabling trading. Though not tested in the first stage, they will be used when carrying out trading tests in the second stage of the pilot.

CREATION OF TOKENS AND CALCULATING TOKEN SUPPLY

During Smart Energy Contract deployment, tokens were created based on the following rules:

- Monthly timestamps start from the beginning of the used data set, which is 01.01.2017
- Time stamps end with the end date of the used data set, which is 31.12.2017
- For each calendar month until the end date of the test, the tokens were created and the Monthly Token Supply was calculated based on the following formula:

$$\text{Monthly Token supply} = \text{Number of days in the month} \times \text{Trading intervals per day} \times \left(\frac{\text{Maximum needed capacity}}{0.1\%} \right)$$

where Maximum needed capacity represents the peak consumption capacity of 2017 & 0.1% represents the minimum potentially tradeable unit

For January 2017, the token calculation would look like:

Capacity: 1483 MW

Days in a Month: 31 days

$$\text{January Token supply} = 31 \times 24 \times \left(\frac{1483}{0.1\%} \right) = 1,103,352,000$$

Calculating every month's token supply accordingly, the Total Token Supply would result in 12,991,080,000 tokens.

The Test

WePower had several hypotheses before the test:

- Median wait time of each transaction is sufficient for the WePower use case
- During certain time periods, it is remarkably cheaper to upload the data
- Using safelow gas price strategy as the minimum of our set gas price will guarantee transaction confirmation in 30 minutes
- We can upload 1 year worth of data into blockchain faster than 1 year.

Regarding the test procedure itself, there were 2 main aspects that we tried to distinguish. These were optimization of cost and optimization of time. This gave us the baseline to optimize the performance over both aspects for supporting WePower's business needs.

Before the test, we had several runs on the local machine to act out the logic of the test and understand the cost and time indications a bit better. Based on that, we budgeted the entire test and clustered different stages.

Uploading the 3,837 corporate energy buyers and the consumption data took us 434 transactions. We started by adding consumption in batches of 15 transactions, after which batches of 3 were used for the rest of the test. All together the total gas used was 1,129,171,462. The average gas price was 9.65921659 Gwei and it cost us 9.366401377 ETH. The average transaction cost for the ETH Main network batch was 2,601,778 Gwei. The test was carried out between 11.10.2018 21:25 and 13.10.2018 13:07.

The transactions can be seen at:


<https://etherscan.io/txs?a=0x9d81027802aa8110540340bb3c552923a131174f>

Total Transactions	434
Gas Used	1,129,171,462
Average gas price (Gwei):	9.65921659
Median gas price (Gwei):	2.5
Cheapest gas price (Gwei):	1.5
Highest gas price (Gwei):	60
Average wait time (s):	1,456.5
Median wait time (s):	203.6
Shortest wait time (s):	0.8
Longest wait time (s):	44,882.1

Table 1: Stats of the ETH main network test

Table 1 provides an overview of the number of transactions and gas used, together with gas price and transaction confirmation time statistics. In Figure 2 we show a more detailed breakdown of various gas prices used and the number of transactions carried out with those prices.

Figure 2: Count of transactions at different gas price



Though the ETH Gas Station states, in regards to their transaction times and gas estimates, that “Estimates not valid when multiple transactions are batched from the same address or for transactions sent to addresses with many (e.g. > 100) pending transactions”, for the first batch of transactions for our ETH Main network test we uploaded data in 15 transaction batches, which we considered to be a reasonably small amount. Those 15 transactions were uploaded with a certain gas price, which was the market price at that moment.

The problem we soon encountered was that gas price levels increased while the transactions were still waiting to be confirmed, meaning that other transactions from the batch had to wait in line until the gas price would again meet our set price level. It meant that at one point we waited roughly 12 hours for our transaction to go through. You can see this in Figure 3, where at 4.1 Gwei took remarkably more time on average for a transaction to go through. Also in Figures 6 and 7 it can be seen that on 12.10.2018 the hours between 7 and 19 are missing for the same reason - no transactions were carried out while the next in line transaction of the batch was waiting.

After those transactions, we changed the logic and started to upload the transactions in batches of 3, which gave us a remarkably faster reaction time in case the gas price changed. After every 3 transactions, the market price was checked and modified according to the strategy chosen. For checking the most up to date gas price levels before we started loading each new batch, our script utilized ethgasstation's ETHgasAPI².

Figure 3: Median time to transaction confirmation per gas price




Figure 4 shows the distribution of transactions based on a chosen gas price strategy. For most of the transactions, we used cost-aware conservative gas price strategies like safelow and average. Where safelow is the lowest price accepted by at least 5% of the miners (and at least 50 transactions have been mined in the last 24 hours at this price), average reflects the price accepted by top miners who account for at least 50% of the blocks mined whereas the fastest price level is accepted by all top miners (estimated over the last two days). To observe the effect of increasing gas price above fast strategy we used gas price 60 Gwei, which even exceeded the fastest gas price strategy levels at the time of testing.

² <https://ethgasstation.info/json/ethgasAPI.json>

Figure 4: Distribution of transactions per gas price strategy in ETH Main network




Figure 5 shows an overview of how transaction confirmation times were distributed proportionally. The majority of transactions were confirmed (included in a block) in less than 5 minutes and more than 25% of transactions confirmed within 60 seconds. The reason behind the rather large proportion of transactions taking more than 30 minutes to confirm was the fact that posted transactions were processed sequentially, so we had to wait for the first transactions to finish, regardless of the higher set gas price in following transactions.

Figure 5: Transaction confirmation time distribution

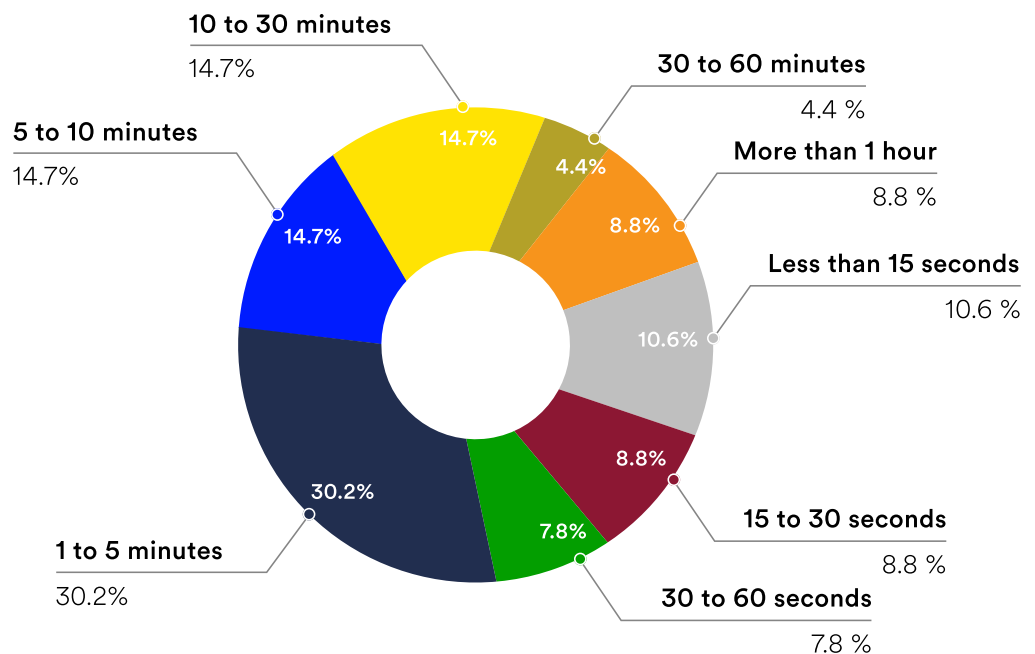


Figure 6 shows the hourly statistics of the average amount of blocks and time we had to wait until our transactions were written into a block. The results generally confirm the assumption that shorter confirmation times can be achieved by increasing the gas price. However, performance depends on other activities being performed at the same time on the blockchain, meaning that increasing gas price does not always guarantee timely confirmation of transactions. For the 3rd until the 10th hour of the test (21 till 07) the uploading logic of batches per 15 was used, this caused a situation where waiting time before our transactions were included in blocks was very high. The count of transactions stayed on the same level and therefore there were no remarkable changes in the gas cost.

Figure 6: Average count of blocks waited and time to block at different hours

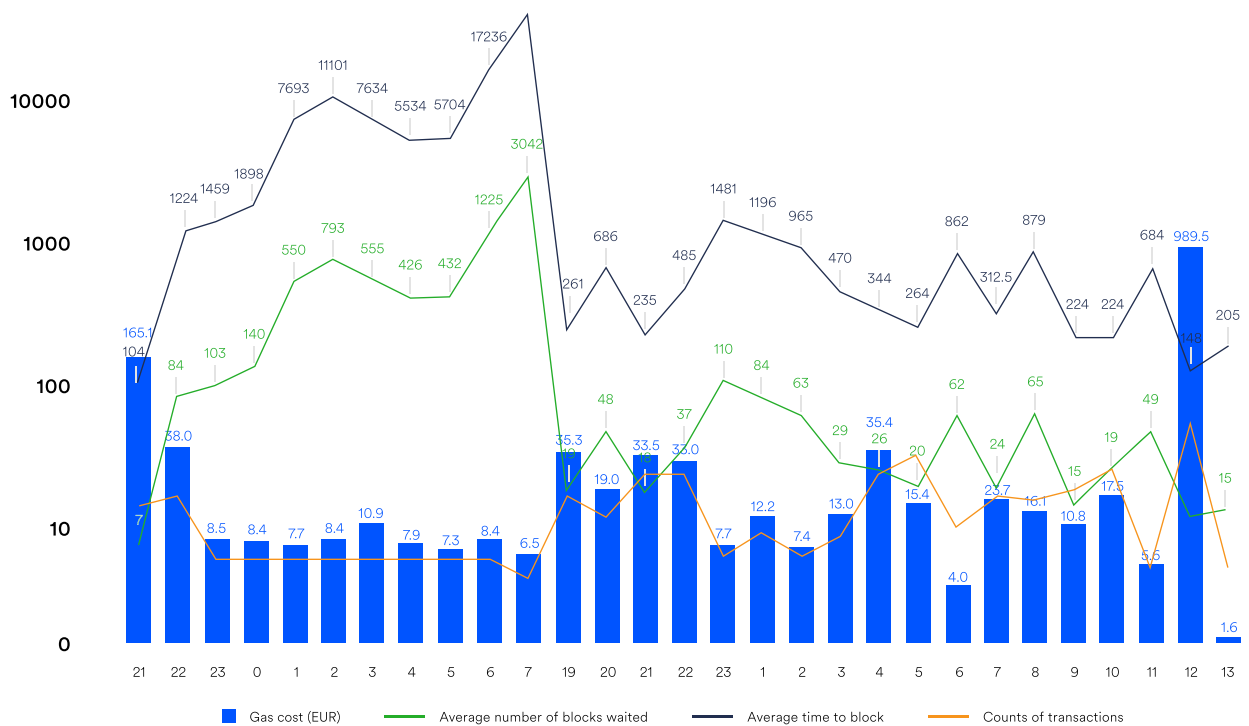



Table 2 highlights hourly minimum/maximum gas prices and transaction confirmation times during “add consumption” transactions. The results support the need for dynamic gas price modification between transactions. Comparing hours 6 and 10 below, both minimum and maximum confirmation times were remarkably higher during hour 6 than the hour 10. While hours 11 and 13 had similar minimum confirmation times at the same gas price level, their maximum confirmation times varied greatly. The average hourly quantity of added consumption data was 4604 data points.

Table 2: Stats regarding the consumption data uploading

Hour	Minimum gas price in Gwei	Maximum gas price in Gwei	Minimum time to confirm in seconds	Maximum time to confirm in seconds
4	1.8	3	20.6	808.4
5	1.7	2	7.5	849.4
6	1.7	1.8	77.9	3914.8
7	2.1	2.4	4.8	812.5
8	1.9	2	32.4	2884.5
9	1.7	2	2.4	1165.8
10	1.5	1.8	3.0	796.6
11	1.5	1.5	146.7	1049.0
12	1.5	60	0.9	981.6
13	1.5	1.5	175.8	230.6


In Figure 7 it is outlined that in situations where the count of transactions in blocks was low, the block size didn't have all the same peaks compared with the transaction count, though the overall correlation can be seen. That means that in general, the count of transactions in the block and the block size have a correlation but we cannot say that this is always true.

Figure 7: Hourly minimums and maximums of total transactions per block and block size



In figure 8 it can be seen that price strategies of average and safelow are very similar and correlate a lot. Also from that, it can be stated that a strong majority is choosing between those strategies. We could assume that for getting your transactions in the block faster than average, a good strategy would be to set a gas price just over the current average, but it would still give a good cost saving compared to fast and fastest levels. From the sample we have, it is not possible to find a clear correlation between the movement of the safelow and the fast and fastest prices.

Figure 8: Loading optimization - hourly minimums of safelow, average, fast and fastest gas prices





Takeaways

Ethereum is mature enough to accommodate contracts with multi-year terms. There are promising alternatives available as well in the development phase, but their infancy is a concern in relation to typical durations of energy contracts.

Ethereum performance indicators are sufficient for WePower to operate the renewable energy procurement and trading platform at this stage. The hybrid platform architecture can be operated with the required cost-performance for the current business needs. In order to meet increasing performance requirements in the future though, WePower will have to consider validating other blockchain/ DLT solutions to identify more efficient systems to run the larger scale trading platform.



Conclusion

The first phase of the Elering tokenization pilot helped WePower to validate and verify the logic and processes that will be at the core of the WePower platform. We have been able to generate Smart Energy Contracts and Smart Energy Tokens based on real-life data on various granularity levels.

During this process of data preparation, we gained valuable experience in dealing with an Energy System Operator. It is possible to exchange data in various formats, but the existence of the Estfeed data hub that provides data in a unified format makes it possible to operate the WePower energy contracting and trading model more efficiently and will be the focal point of further innovation when more data hubs are implemented across markets in various countries.

We acknowledge that the scalability of Ethereum blockchain currently has limitations. However, the problem is being tackled by Ethereum developers with plans to implement sharding, i.e. partitioning data into subsets, and moving from energy-intensive Proof of Work to a more environmentally friendly Proof of Stake consensus model. Ethereum co-founder Vitalik Buterin has stressed that there is a need for solutions that mitigate expensive storage costs that could escalate exponentially as the system expands³.

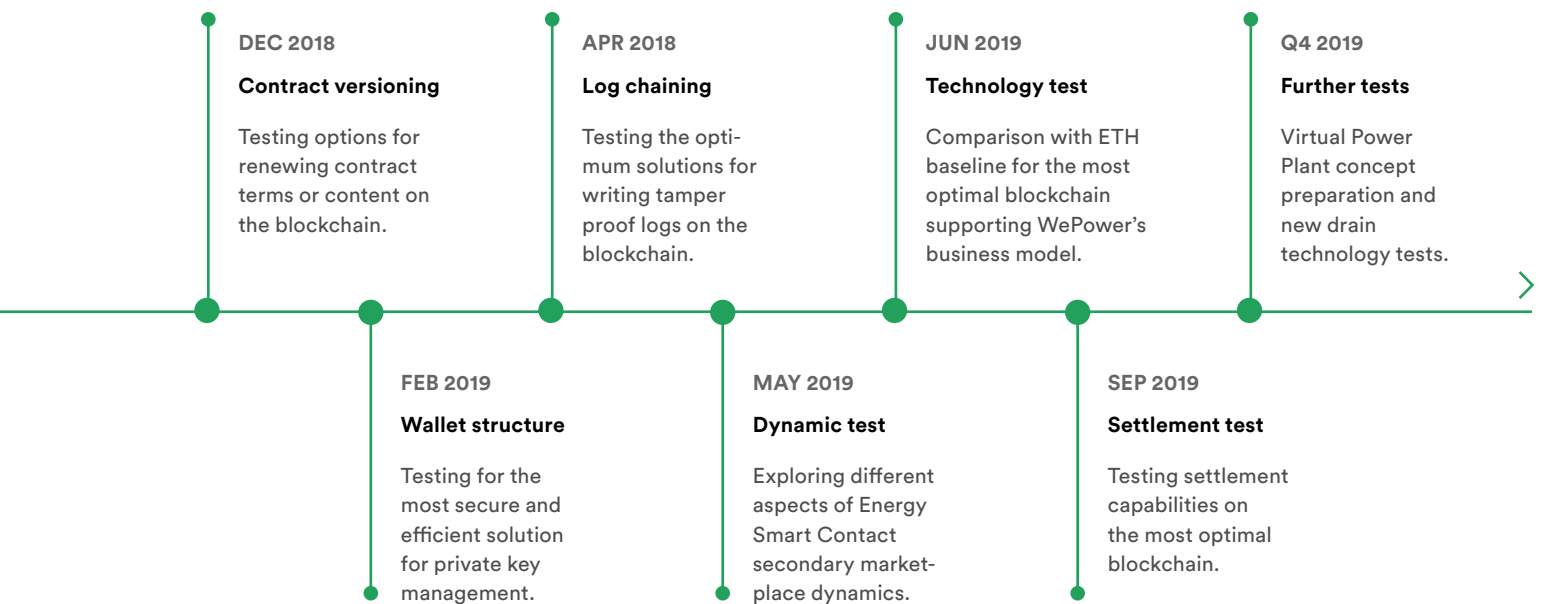
³ <https://www.coindesk.com/modest-proposal-vitalik-unveils-multi-year-vision-ethereum/>

The balance between cost and transaction confirmation times has to be considered continuously while communicating with the blockchain. Practical lessons from coding perspective will be analyzed and implemented in the coming stages of WePower’s pipeline. While Ethereum is currently one of the most mature blockchain solutions, supporting smart contracts, a fully decentralized application for large-scale autonomous usage is not yet feasible for large-scale energy trading on the blockchain. For now, a hybrid solution will be utilized while monitoring the rapid development of Ethereum as well as other blockchain/DLT technologies.

Next Steps

The first stage of the Elering pilot was just the beginning of WePower’s R&D activities.

Figure 9: Roadmap for the upcoming research subjects



OVERALL VISION

For the next steps our R&D is focused on the features that are required to further optimize WePower's platform for large scale commodity trading purposes. The roadmap is leading the way to concepts that support WePower in growing to a fully functioning virtual utility. For that WePower has defined a series of features and use cases that need to be covered at a detailed level. For supporting WePower's primary value offer - contract versioning, wallet structures, log chaining and real electricity data and price settlement on blockchain have to be thoroughly analysed and tested.

FEATURE OPTIMIZATION

We are going to analyse and test how we can optimize changing contracts through versioning. This will help WePower optimize the way that smart contracts have to act on the code and process level, when contract holders start the activities that need modifications in the contract content.

For writing immutable transaction logs into blockchain we are going to analyze different methods of how we can optimize the amount of data written into the blockchain - at the same time being sure the logs remain tamper proof.

MARKET DYNAMICS

When we have the essentials covered, we will look deeper into energy market behaviour. This is also known as the second phase of the Elering pilot, where we will explore different aspects of secondary Smart Energy Contract marketplace dynamics. Meaning that we will go into more detail regarding various events that can happen with market participants and contracts during the Smart Energy Contract lifecycle (including peer-to-peer trading of contracts, change in consumption profiles, settlement and netting of energy production data with electricity prices etc). This serves as preparation for concerns that might come up in the various stages of WePower's business processes from the blockchain view.

SECURITY

Ensuring cybersecurity is correctly addressed throughout the entire project lifecycle is a critical success factor. With additional functionality being added to Smart Energy Contracts and the Platform's functionality, we will continue to ensure these functions do not compromise the security of users.

WePower will prepare a paper based on the main observations from our own activities and domains' best practices.

TECHNOLOGY CHOICE

Once we have finalized these stages we will take a critical look at the different available technologies for WePower to use, as the domain is maturing day by day. Given that we will have done many tests on Ethereum by that point, we will have a good baseline to compare it to other blockchains. This serves the purpose of WePower not wasting time and resources developing functions on the blockchain that are not feasible for the future. Settlement tests will happen on the best possible blockchain for WePower's needs. During the settlement tests, we will examine the possibilities of dealing with the settlement functionalities on the blockchain.

After that, we will dive into more conceptual analysis, which will need various proof-of-concept tests before we can choose a feasible track moving forward.



