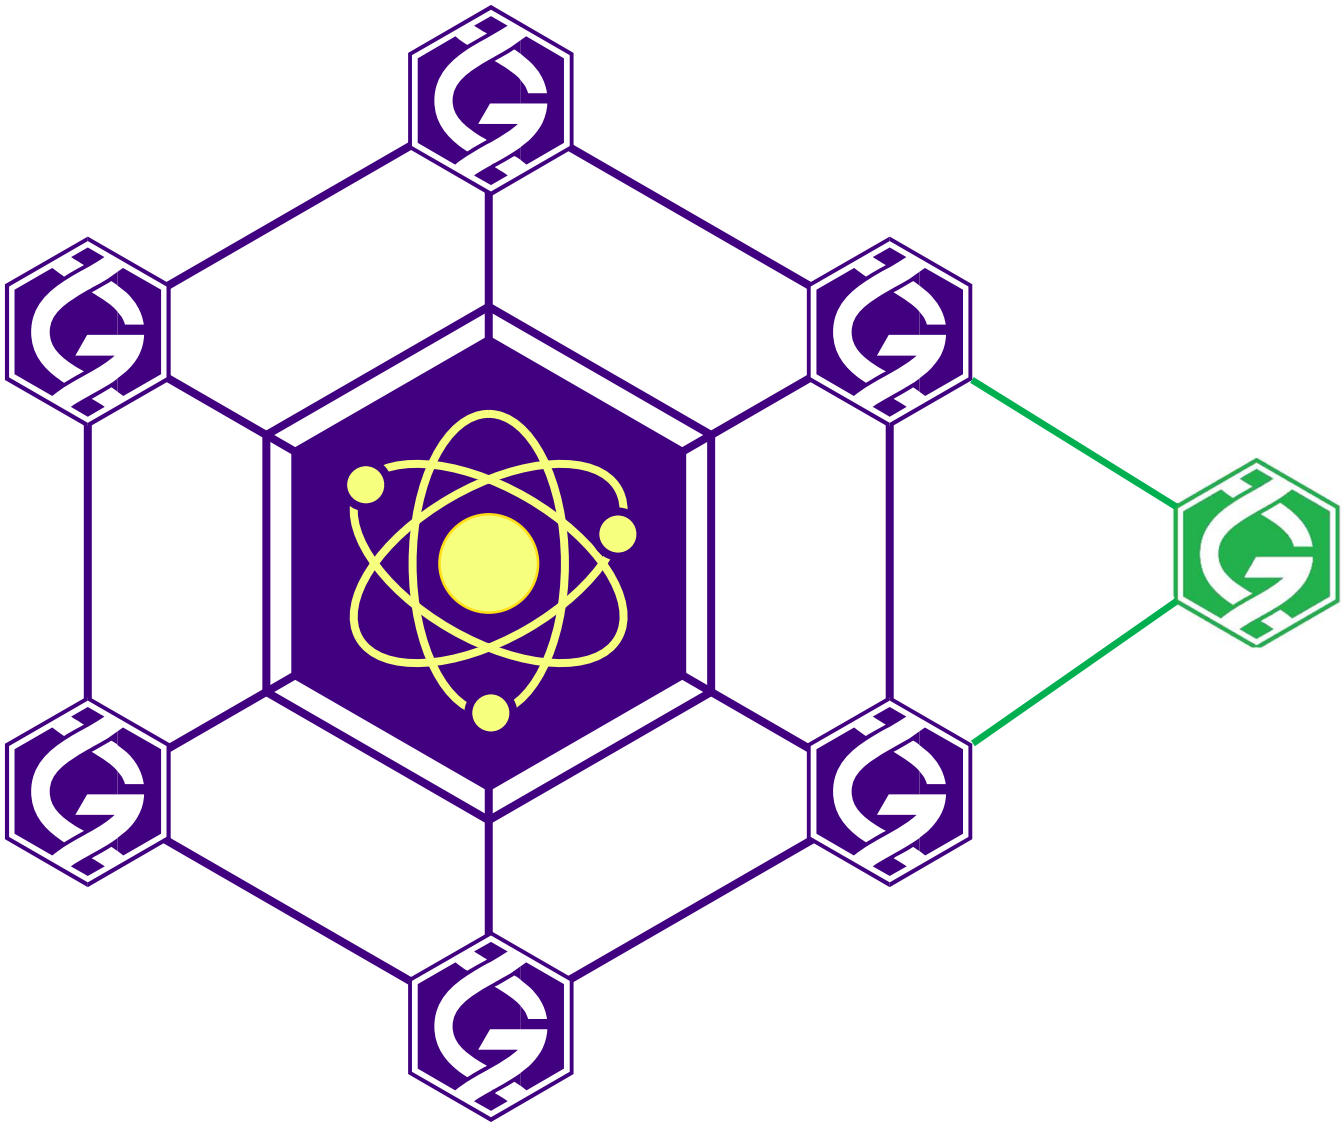

# GRIDCOIN GREEN PAPER

## A PROPOSAL FOR A CLOSED LOOP ECONOMIC SYSTEM SERVING SCIENCE WITH CAPITALISM

VERSION 0.0.1 [BETA]

Gridcoin

# Table of Contents

# Objective

The purpose of this proposal is to engineer an economic pathway for massively scaling the Gridcoin network. In brief, this involves "greenlisting" computing projects from for-profit companies without awarding them minted GRC.

# Disclaimers

There are several criticisms of this idea. I have done my best to anticipate and answer them throughout this paper. I look forward to discussing more with the community. Here are some key things to keep in mind:

1. Nearly all the numbers presented should be thought of as rough approximations. They are accurate to the best of my ability, but many are based on constantly fluctuating measurements (peta-FLOPS of the BOINC network, the price in dollars of Ethereum, etc.). I believe my system is still functional so long as they are within one to three orders of magnitude.
2. Additionally, measurements of total computing power are fraught with difficulties and inaccuracies. Peta-FLOPS are a reasonable approximation but must be understood in light of some serious caveats. Differences in hardware, algorithms, implementations, and even compiler tuning can have a large effect on the efficiency of computation.
3. Much of this paper will discuss prices of cryptocurrencies in dollars, including GRC, and this is at odds with the vision of many leading voices in this community. **I wholeheartedly agree that a myopic focus on price distracts from the truly inspiring things that could be built on an Open Economic Network devoted to science. Those inspiring dreams must continue**. However, a common refrain is "we hope crunchers will come because of greed and stay for the science." This specific proposal is about a way to make the GRC currency profitable in dollar terms for the express purpose of siphoning profit motivated Ethereum miners. That is a prerequisite to a significant expansion of the network and would be an enormous benefit to the BOINC projects we support.
4. From Wikipedia: "A green paper is a tentative government report and consultation document of policy proposals for debate and discussion. A green paper represents the best that the government can propose on the given issue, but, remaining uncommitted, it is able without loss of face to leave its final decision open until it has been able to consider the public reaction to it." I am sure the Gridcoin community will uncover serious, and possibly fatal, flaws with this idea. If I can merely inspire better ideas, then this paper will have been a success.

# Problem Statements

1. Gridcoin currently trades for about $0.01. It is worth far less than the electricity required to mine the coin. Until that changes, no profit driven miner will switch to crunching BOINC.
2. Gridcoin developers are brilliant and grossly undercompensated. Any proposed solution must be simple to implement and require minimal changes to the codebase.
3. The Gridcoin supply seems to be held primarily by several mostly inactive whales. As they hold a significant amount of voting weight, forking the blockchain and "starting over" is considered off-the-table for this proposal. A solution that leaves the whales untouched is preferable.
4. Complicating this, at least one of the whales may have begun to slowly sell a significant amount of GRC. This has put downward pressure on the price. Any solution must be able to soak up an enormous supply of GRC for demand to finally outpace supply and raise the price.

# Some Esoteric Economics Philosophy

[The following is my personal opinion and understanding and should not be taken to represent the opinions of the Gridcoin community at large]. Most cryptocurrencies have a chicken and egg problem. Bitcoin has many properties which make it a useful store of value (scarcity, predictability, divisibility, etc.). However, none of these makes Bitcoin inherently valuable in the way that land is inherently valuable. Bitcoin's value in real, tangible, even dollar denominated terms is derived from the OEN's shared belief in its future value. Bluntly, people are willing to trade goods and services for Bitcoin now because they believe that in the future other people will be willing to trade even more goods and services for those same Bitcoins assuming broad economic growth and a known limited supply of Bitcoin. This chain of network belief works well once established but is difficult to get going. At some point someone must be willing to give up two pizzas for thousands of Bitcoins.

This method of starting a cryptocurrency value chain seems to have a significant first mover advantage. Thus, I prefer to assume that such a value chain is unlikely to ever boost the value of Gridcoin. Instead, GRC's value must be derived from something other than accidents of speculation. Gridcoin is not an investment. One method is to build interesting scientific services on the Gridcoin blockchain which can only be accessed in exchange for GRC (see Disclaimer #3). Examples include scientific paper publication, scientific accreditation, and network governance (voting). A parallel and complimentary method is to inject value into the system by linking the Gridcoin network to existing economic markets, namely the cloud computing industry. One potential implementation of this concept will be described here.

# Some "Napkin Math"

Here are some key numbers:

The price that GRC would need to reach ($/GRC) for mining with a GTX 1070 to:

- Break even with electricity costs ($0.10/kWh) = $0.03
- Beat out the least profitable alt-coin on WhatToMine.com (Masari) = $0.04
- Become more profitable to mine than Ethereum = $0.29

Minor caveat – if everyone currently crunching the whitelisted projects simply setup Gridcoin and began getting paid, then competition for magnitude would spike, causing these break-even prices to roughly triple. I view this as unlikely given my research into common attitudes towards profit-seeking in the BOINC community, and so is not considered here.

- Very approximate FLOPS / Ethereum Hash = 250,000
- Hashrate of the Ethereum network at 2018 peak (prior to Eth ASICs) = 73 Exa-FLOPS
- Current Total BOINC power = 27.9 Peta-FLOPS

# The Proposal

1. Modify the Gridcoin protocol to allow for whitelisted projects to be granted different shares of the same minted GRC, including 0%.
    a. A whitelisted project which receives 0% minted GRC would be considered "greenlisted."
2. Modify the "rainbymagnitude" RPC command to divert 20% of rained GRC to the Foundation Fund, and limit the total GRC that can be rained to no more than 2240 per day per project (1.2 times the amount currently minted to whitelisted projects).
    a. The 20% diverted to the Foundation Fund is referred to here as the **Service Fee**.[1]
    b. The 1.2 multiple funding cap is referred to here as the **Funding Limit**.[2]
3. Allow for-profit projects to be granted greenlist status. These for-profit projects would fund and incentivize crunching on their project using the modified "rainbymagnitude" RPC command.
4. Allow the Foundation Fund to sell GRC to greenlisted projects at a negotiated price well above market rate (likely $0.29/GRC to begin). This is to get around the relatively low liquidity of GRC on exchanges currently.

There are some key things to note with this. First and foremost, it must be reiterated that the objective of this modification is to strengthen the economics of the Gridcoin network for the express purpose of furthering its contributions to science. Any hopeful for-profit company must still go through a community vote to ensure that no project is greenlisted that conflicts with the values of the network as stated in the Gridcoin Whitepaper. As an extreme example I would not support greenlisting a weapons company. On the other hand, there are many for-profit companies that do important work to further science and beneficial technology. If Tesla needed more processing power to hone their self-driving algorithms, I would support adding them to the network.

# An Example With Numbers

The company Open Science Inc. wants to train an open source artificial intelligence algorithm which would enable them to make money while solving world hunger. They have some generous angel investors but need access to a lot of power for about three years. For simplicity, let us assume that:

- The total number of crunchers on the Gridcoin network stays relatively constant at current levels
- We greenlist only one project at a time
- All 15 whitelisted projects continue to be rewarded with equal amounts of GRC (1/15 of total supply ~= 1,867 GRC/day, ~=680k GRC/year)
- The greenlisted project purchases and funds their crunching with an amount equal to that of one whitelisted project, plus 20% Foundation fees (1,867 x 1.2 = 2,240 GRC/day)
- Open Science Inc. wants to purchase all the required GRC (2,452,800) up front from the Gridcoin Foundation at a negotiated price of $0.29/GRC in anticipation of rising GRC prices and to get around the low liquidity of GRC on SouthXchange.

[1] The Service Fee is likened to the fees paid to companies like Uber and Airbnb for the privilege of using their technology platform to connect buyers and sellers. 20% is roughly equivalent to the fees charged by these two companies.
[2] The Funding Limit is in place to protect whitelisted scientific projects from being overwhelmed by deep-pocketed greenlisted companies. 1.2 ensures even competition and accounts for the 20% funding fee. This could be adjusted as the total power of the network scales.

Using these starting assumptions, it would cost a business $711,312 to rent 364 Tera-FLOPS of computing power for three years. The cheapest alternative would be to buy the hardware and build an equivalent datacenter at a cost of around $12,000,000.[3] The following chart demonstrates the value proposition of the Gridcoin network.

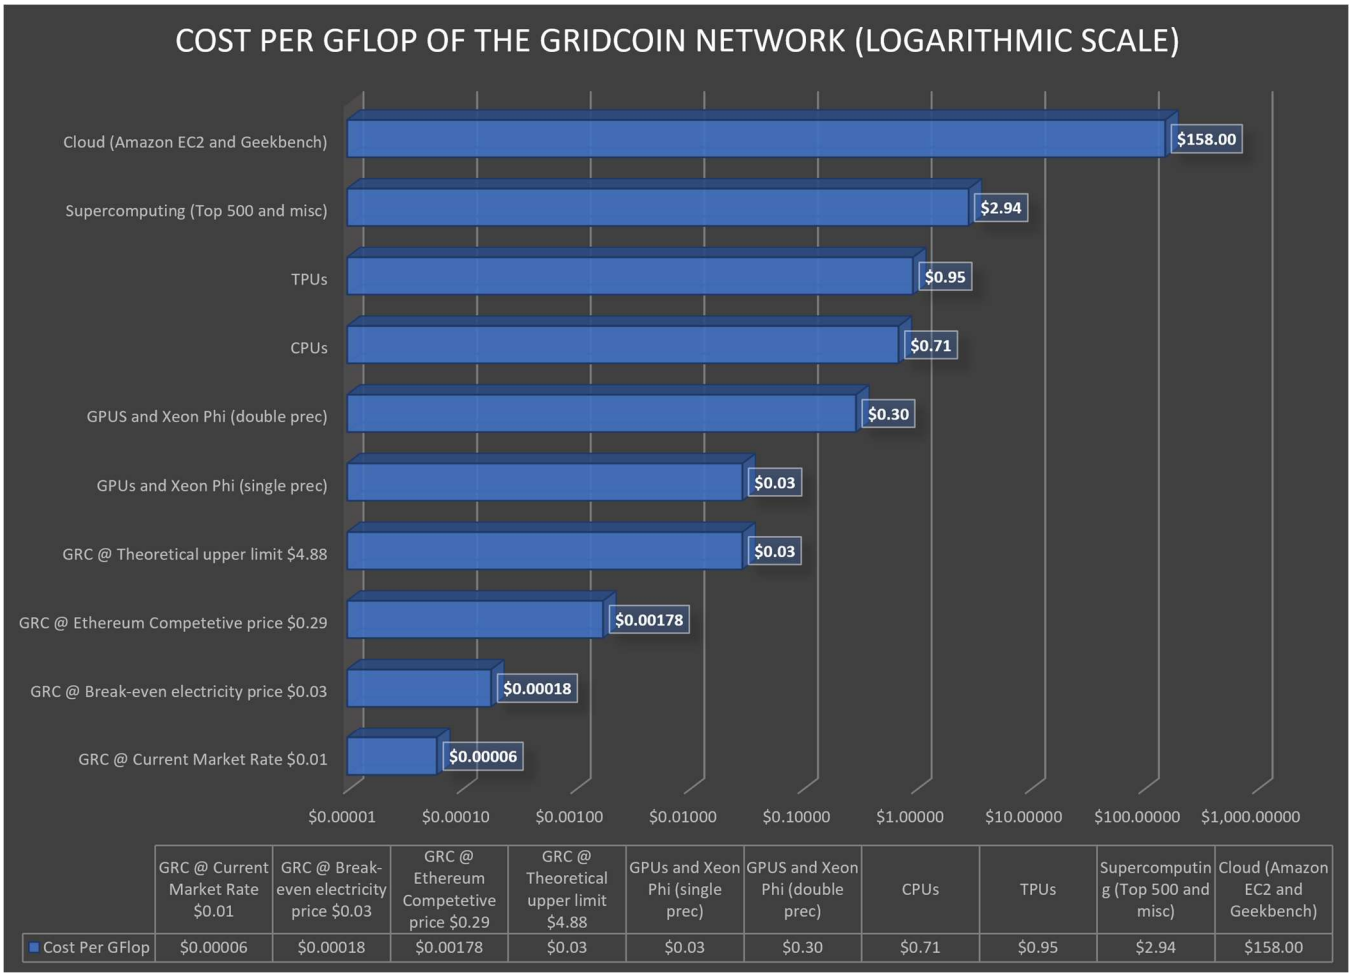


| | GRC @ Current Market Rate $0.01 | GRC @ Break-even electricity price $0.03 | GRC @ Ethereum Competetive price $0.29 | GRC @ Theoretical upper limit $4.88 | GPUs and Xeon Phi (single prec) | GPUS and Xeon Phi (double prec) | CPUs | TPUs | Supercomputing (Top 500 and misc) | Cloud (Amazon EC2 and Geekbench) |
|---|---|---|---|---|---|---|---|---|---|---|
| Cost Per GFlop | $0.00006 | $0.00018 | $0.00178 | $0.03 | $0.03 | $0.30 | $0.71 | $0.95 | $2.94 | $158.00 |

Source for non-Gridcoin comparisons: https://aiimpacts.org/current-flops-prices/

Note : This source is from 2017, Moore's law may have modified the numbers somewhat recently. I believe they are still close enough for rough order-of-magnitude comparisons.

## Another Comparison

It should be noted that comparisons and calculations based on FLOPS are notoriously inaccurate. Jim Owens wisely stated once that "a FLOP is not a FLOP" (see Disclaimer #2). To add weight to the argument, I ran a comparison with a similar company, Kings Distributed Systems. I used the makeup of the Einstein@Home project to approximate the total number of CPUs and GPUs attached to each whitelisted project. This came to 8,428 CPU cores and 567 GPUs per project. It would cost $1.5M per year to rent this many CPU cores and GPUs from Kings DS, or 6.5 times the cost of Gridcoin from our hypothetical example above.

[3] Based on my own calculations and assumptions. This number is in close agreement with that posted on the BOINC Wiki: https://boinc.berkeley.edu/trac/wiki/BoincOverview

## How can this be so cheap?

I believe the numbers prove that Gridcoin can take advantage of inefficiencies in the market in the same way that other actors in the sharing economy have found significant market opportunities. The most readily comparable examples are Uber and Airbnb. Both companies are essentially clever technology platforms that enable the economic mobilization of previously underutilized capital resources, namely personally owned vehicles and spare bedrooms. Likewise, Gridcoin is a brilliant technology that enables the economic mobilization of underutilized personally owned processing power.

## My Favorite Takeaway from this data:

Ethereum miners are underpaid.

# Discussion of Effects on The Network

This proposal as executed in the simple example above would provide several critical things for growing the strength of the Gridcoin network. First, it would provide the Foundation with a sizable amount of dollars with which to pay developers ($0.7M). It would also provide a steady income stream of Gridcoin with which to refill the Foundation Fund, approximately 163,000 GRC per year. If additional greenlisted projects bought GRC on the open market, demand for the currency would go up considerably and may eventually soak up the excess supply currently being dumped by hopeless whales.

*An additional benefit of a more profitable value for GRC is that it would allow a reasonably good approximation for the size of the Gridcoin Foundation's donations to science in dollar terms. At $0.29/GRC, this is $1M worth of computing power donated per project. Being able to attach sizable dollar figures adds weight and credibility to the contributions of the network and would benefit Gridcoin's standing in the scientific community. These figures would be further validated by the fact that for-profit companies would be paying the same amount of money for an equal amount of processing power. Finally, we could argue for acknowledgements in scientific papers produced by whitelisted projects: "This project was made possible by a grant of $1M from the Gridcoin Foundation."*

Eventually, as demand rises to meet supply the price of GRC would rise, first to break even with electricity costs, and later (hopefully) to compete with Ethereum. Best case is that this happens around the same time that Ethereum transitions to Proof-of-Steak and 73 Exa-FLOPS of GPUs are suddenly out of a job. After this, the network could scale arbitrarily to meet market demands for distributed computing power.

## Protecting Science

As stated at the outset, the purpose of this proposal is to further the Gridcoin network's ability to benefit science by capturing and diverting market forces. It is important to ensure that as additional projects are greenlisted they do not divert significant amounts of power away from whitelisted projects. This is the reason for the Funding Limit previously discussed. This would be less of an issue as GRC becomes more profitable to mine and attracts additional profit-motivated crunchers. As the total power of the network scales up, the Funding Limit and the number of greenlisted projects should also be allowed to scale up to increase the strength of the Gridcoin economy.

It is also important to remember that whitelisted Work Units are, in most cases, a limited resource (with the possible exception of Collatz Conjecture). The Ethereum GPUs are, collectively, 8,242 times more powerful than all of Gridcoin. The goal should be to make available to science as much computing power as is needed, and no more. The rest would, ideally, remain connected to the Gridcoin network on greenlisted projects until more science is needed.

## Scalable Minted Rewards

The reason for Proposal #1 – Scalable Minted Rewards is to allow for some project's minted rewards to be 0%, i.e. greenlisted. An interesting side benefit is that whitelisted projects could have their rewards scaled according to their work units available. This would enable small projects with relatively limited work units and long time scales to still participate in the Gridcoin network in a whitelisted status. This would be especially important as the total power of the network scales up to prevent smaller projects from running dry and risking becoming greylisted.

## A Closed Economic System

Adding greenlisted projects would finally realize a closed economic loop for Gridcoin. GRC is bought by greenlisted companies, who pay it to crunchers for their computing power. Crunchers can in turn sell it to greenlisted companies. Minted GRC would continue to incentive whitelisted scientific projects and would serve as an anchor for the GRC/Computing Power market dynamics (companies would have to closely match their payments to the amounts minted to whitelisted projects to stay competitive). The Foundation would benefit from both the rise in $/GRC prices and the associated Funding Fees, which they could use to pay for increased development on the platform.

## A Note on Cheating

For-profit companies that have invested a significant amount of money have a deeply vested interest in preventing and policing cheating. If people can cheat, the company will in effect be getting less computation done for the money they paid. Cheating is not expected to be a significant issue on green-listed projects.

# Prerequisites and Implementation

The next step for this proposal is to allow the community some time to discuss, critique, and modify it. After that, if the proposal looks promising it should be put to a vote via Gridcoin poll to ensure broad support.

I have done my best to come up with a plan that would require as little development effort as possible. However, I must defer to the core developers to comment on feasibility of the two primary changes discussed in Proposal #1 & #2.

The bulk of the implementation work would be administrative in nature, including:

1. Forming an official Gridcoin 501c3 Nonprofit or LLC.
2. Creating business communication materials such as flyers, presentations, and websites.
3. Advertising and reaching out to interested businesses.
4. Greenlisting the first project as a trial run.

# Possible First Candidate

I have thought of one possible company that may be promising as an initial candidate. I am sure the Gridcoin community will think of many more. The Numerai project is a hedge fund which incentivizes data scientists to create machine learning models that can predict movements in the stock market. Competitors are paid in the Numeraire (NMR) ERC20 token. Anecdotally, the very first winner of the weekly competition (several years ago) was paid about $100 for his winning contribution. He laughed and said "I spent about 20 times that much on the compute to generate that model!" The Numerai community is intimately familiar with cryptocurrencies, computer science, and cloud computing, and could bring many brilliant minds into the Gridcoin space. Their machine learning models may be able to benefit from increased computing power. Implementation would be difficult since there are hundreds of data scientists training and competing different algorithms. The Numerai company or community would have to setup a centralized BOINC server which would accept work units from the data scientists and distribute them to the Gridcoin network. They recently released their updated Numerai-Compute framework for automating certain associated aspects on Amazon Web Services, but they have expressed interest in extending this to the Golum distributed computing project as well.

# Possible Second Candidate

Likely further afield, the prospect of greenlisting a power company would have even more interesting use cases. In 2019, a large utility company in France, EDF, began working with iExec, a "decentralized marketplace for cloud resources" which uses RLC as its native token.[4] What this demonstrates is that large utility companies may have significant requirements for computing power. A greenlisted utility company may be willing to take GRC as payment for electricity consumption, adding a fascinating layer to the closed economic loops within the Gridcoin economy.

# Downsides to Gridcoin

In its current implementation, Gridcoin is most easily applied to projects built on the standard BOINC infrastructure. This is by no means a requirement, and could change soon, however BOINC was not designed with much consideration for data security and privacy. A company with even the slightest desire to ensure the security of proprietary business data and algorithms will likely not choose Gridcoin. Marketing and advertising would have to focus on companies friendly to Open-Source principles.

The Gridcoin system works best on large, long-term computation projects (greater than one month). The process of polling and greenlisting as well as BOINC's RAC "charge up" formulas would make operating shorter-term projects difficult. This would further limit the pool of potential companies.

Gridcoin is open-source and could be easily duplicated without the inflationary mint boosting scientific projects. A network identical to Gridcoin without the mint would be cheaper to rent than Gridcoin. Gridcoin's only advantage in this case is it's first-mover status and sizable pre-existing network.

[4] Utility Giant EDF Deploys Blockchain Application on Ethereum and iExec

# Summary

Purchasing computing power on the Gridcoin network is several orders of magnitude cheaper than any alternative. Selling a portion of the network's computational power to respectable for-profit companies would re-fund the Foundation, increase and stabilize the value of the GRC currency, and provide closed economic flows to the Gridcoin economy. In the long term, this may facilitate the capture of a large portion of the Ethereum network's GPU infrastructure…for science.

## Closing Remarks

Many thanks to J-Ringo, Jim Owens, RoboticMind, ILikeChocolate, and the many other Gridcoiners that encouraged me to dive in and contribute some ideas.

If you thoroughly enjoyed this paper and simply must tip the author, here is my wallet address:

S1fSQwAJD5bzBCucsCRs6kPUCwhfGAEEFt

But most of all, I hope you generate your own better ideas and share those back. Sincerely,

-AgentP