

The logo for DASH, featuring the word "DASH" in a bold, blue, sans-serif font. The letter 'D' is stylized with a horizontal bar extending to the left, creating a sense of motion or speed. The background is a light gray gradient with a large, white, rounded rectangular shape on the right side.

## **Dash Governance System**

Dash Core Team Response to the report "*Dash Governance System: Analysis and Suggestions for Improvement*"

Ryan Taylor – [ryan@dash.org](mailto:ryan@dash.org)

[www.dash.org](http://www.dash.org)

Note: the purpose to this article is to provide an official response and clarification to the report "*Dash Governance System: Analysis and Suggestions for Improvement*" published by "The Veritas Team" - Dmytro Kaidalov, Andrii Nastenکو, Oleksiy Shevtsov, Mariia Rodinko, Lyudmila Kovalchuk, and Roman Oliynykov on <https://iohk.io/research/library/#NSJ554WR>

## Table of Contents

Acknowledgement.....	3
Introduction.....	3
Ability to Adapt.....	4
Minor Corrections and Clarifications.....	4
Weaknesses.....	5
4.1. Superblock Validation Attack .....	5
4.2. Multiple Voting.....	6
4.3. Discarded Votes.....	6
4.4. Masternode Voting Incentives .....	6
4.5. Voting Representation.....	6
4.6. Voting Delegation .....	7
4.7. Decreasing Treasury Fund .....	7
4.8. Flexibility of Treasury Monetary Policy .....	7
4.9. Code Review .....	8
Proposed Solutions.....	8
6.2.1. Flexible Participation Requirements .....	8
6.3. Implementation of Delegation with Liquid Democracy .....	9
6.4. Assigning Categories to Each Funding Proposal.....	9
6.5 Introduction of Paid Expert Roles.....	9
6.6 Expanding List of Voting Entity Types.....	10
Other Weaknesses.....	11
Summary.....	11

## Acknowledgement

We wanted to recognize and thank the Veritas Team for its professional assessment. The depth and quality of their assessment clearly required significant investment of time and talent both to sufficiently understand the functioning of the system, but also to craft the report content. The Dash core team is pleased they found Dash's governance system to be of sufficient sophistication to warrant such research.

## Introduction

The Dash core team deeply appreciates the thorough review of the Dash Governance System and unreleased 12.1 code performed by The Veritas Team (<https://iohk.io/research/library/#NSJ554WR>), hereafter referred to as the "Veritas Report". Its contents can be helpful in improving the system and helps facilitate discussion on the merits of various governance approaches and systems.

Indeed, as Dash gains practical experience with its own governance system, situations have revealed flaws and shortcomings that the team plans to resolve, though the urgency, impact, and importance of each resolution varies. The need for ongoing improvement and greater flexibility was a major impetus and driver for many of the capabilities deployed in 12.1. Part of the benefit of migrating the budget system to the architecture of 12.1 is to allow greater flexibility to deploy improvements to the system in an object-oriented fashion. This both enables a greater range of functionality from the system and enables faster deployment of improvements without the need to change network protocols. The resulting system enables Dash to iteratively improve the governance system going forward and systematically address shortcomings.

The Dash core team found the Veritas Report to contain a high-quality description of the governance system that will initially deploy to mainnet with the deployment of 12.1. This initial deployment intentionally mimics the capabilities of the 12.0 governance system and therefore inherits many of its shortcomings. However, with the new architecture, Dash intends to address many of these shortcomings in a completely decentralized way.

There are two reasons for mimicking the capabilities of 12.0. First, it reduces the risk to the network by changing only the architecture and avoiding concurrent changes to functionality. Second, it ensures that iterative changes to governance are approved and implemented by the network itself, rather than through "forced decree" from the core team. By starting with existing functionality and iterating from there through network-approved changes, these issues are avoided.

The Veritas Team was also helpful in identifying one potential attack vector, which was reported to the team and resolved prior to the issuance of the Veritas Report.

We also agree with many of the report's findings. For example, it is apparent that the current model is unable to scale. If Dash were to match Bitcoin's market cap (and our goal is to exceed it), the budget system would currently provide over \$100 million in project funding annually. It would be infeasible for masternode owners to properly evaluate spending proposals at the existing level of granularity at such a scale, which would likely consist of hundreds of proposals every month.

There are also a number of criticisms or suggested improvements with which we disagree, either because we consider them to be well-considered and valid design choices rather than flaws, because the Veritas Report failed to evaluate the full trade-offs associated with their proposed solutions, or because the underlying evaluation of how the system works was inaccurate (which we acknowledge is understandable with such a complex system).

Finally, because Dash has the benefit of practical experience with its governance system, we identify several additional flaws that the Veritas Team overlooked. We will use this opportunity to document some additional flaws and potential solutions.

## Ability to Adapt

Unlike many other projects in the digital currency space, it is worth noting that Dash is well-positioned for change because it contemplates clear and explicit mechanisms to adapt and evolve any aspect of its functioning including its protocol, governance system, and incentive structure. Dash is also not dependent on its “contractors” to make decisions on its behalf (we will explain what we mean by this statement). Most first-generation networks are inflexible. They are dependent on miner support to implement protocol change and are therefore trapped in a mode of funding in which 100 percent of block rewards and fees are allocated to one set of network participants and the decision to alter that arrangement resides with the same conflicted and “non-staked” parties. The ability to adapt the entire network in a coordinated fashion may be one of Dash’s greatest strengths.

## Minor Corrections and Clarifications

There are a couple of minor corrections to the description of the system and how it functions.

In paragraph three of section 2.3, the authors claim that voting decisions regarding specific proposals can be made until the date of the payment. While it is technically accurate that votes can be submitted to the network, the budget finalization procedure may be initiated as early as three days prior to the superblock. This means that the effective cutoff for voting may be up to three days prior to the budget payment date.

The information in section 5 is also known to be incorrect and may be misleading.

Firstly, section 5 is an analysis of the Dash-denominated allocations only, and makes no adjustment for the value of Dash at the time the proposal was submitted. For example, in the fall of 2015 the value of Dash was below three USD, and was worth roughly four times as much in the fall of 2016. This causes expenses from 2015 to constitute a disproportionate weight in the analysis. For a more accurate picture of how money has been spent thus far, a more rigorous analysis is needed which also considers the value of Dash at the time of the proposal.

Secondly, section 5 relies on information sources that are inadequate to determine how funds were actually spent. For example, the “core team salary” budget contains funding for software development, project management, business development, finance, QA, marketing, and a number of other functions, but those costs appear to be allocated wholly to “Software development” in the analysis.

Thirdly, it is misleading to portray the largest applicants to the system as having “received” funding. Other than direct cost reimbursements, neither Evan Duffield nor Ryan Taylor have received any funding except modest allocations from the core team salary budget (and Evan discontinued any salary after January 2016). However, as founder and director of finance respectively, these individuals submit nearly all proposals on behalf of the core team and on occasion for third parties. Funds associated with these proposals are almost entirely redirected to contractors, third-party vendors, or otherwise used for the benefit of the network. The core team also can act in the role of escrow for third parties wishing to receive funding from the network. For example, Wall of Coins sought funding from the network directly, but the high cost and concerns over recourse for non-delivery caused the core team to act as intermediary, requesting the funds, and distributing the payments in conjunction with sufficient completion of milestones. All financials of the core team,

including balances for each budget, are published on a quarterly basis. These financial reports offer a more accurate depiction of how monies are actually spent from core team proposals.

The most effective networks would necessarily be the ones with a cohesive strategy, cohesive initiatives, and all oars rowing in the same direction, with a minimum of activity directed toward efforts that don't support the strategy. While the Veritas Team expressed concern that funding allocations are primarily going toward the core team because it may lead to inefficient decisions, we actually think the danger of inefficiency is much higher with full decentralization of all aspects of the project, including its primary leadership and decision-making approach. However, any centralized aspects that are created for efficiency or effectiveness reasons need to be created by the masternodes and remain revokable by the masternodes, so that at all times the network decisioning remains fully decentralized in practice. We think it is natural and healthy that most of the funds would be directed in a cohesive manner, which is not necessarily conflicting with a concentrated allocation. We also think that it is natural and healthy for multiple groups of contributors to receive funding from the network. Regardless of our personal beliefs regarding the most effective approach, we believe the network should ultimately decide what is best, rather than asserting either highly centralized allocation or highly decentralized allocation as superior.

## Weaknesses

Here we will respond to each of the asserted weaknesses of the governance system.

Note that there have been multiple versions of the Veritas Report posted to the IOHK.io website, with different section numbering. In the Veritas Report document as it appears upon the issuance of this report, the section numbering is inconsistent between the table of contents and the body of the report. For the purposes of this response, we refer to the numbering as it appears in the body of the report, to allow easier reference when reading the suggested weaknesses and our response. Should the Veritas Team continue to edit their report and change the section numbering, be aware that this section refers to the section entitled "Weaknesses".

### 4.1. Superblock Validation Attack

*>Modeling an attack: ...*

*>At this moment the finalized budget F1 is no longer valid for the superblock S1. All honest nodes will therefore reject S1 and all further blocks.*

This is incorrect, as the security model here is slightly different. Nodes almost never validate past superblocks - they typically only validate "current" superblocks when they are online and synced. If the chain tip changes, it is true that past superblocks could be reevaluated with funding constraints enabled. However, nodes that were simply offline/non-synced when past superblocks were created trust the majority of other nodes to behave honestly and accept such superblocks as is (still verifying regular rules like max subsidy, amount of work, etc).

Note that temporal differences in superblock validity are not on their own sufficient to fork the network. Spatial differences (between nodes on the network) are also needed. While it may be possible for an attacker to induce spatial discrepancies by routing objects or votes preferentially to select nodes, an attacker would need to gain control of a significant share of the masternodes at the time of S1 and figure out object routing schemes that would create spatial differences between nodes in order to fork the network, since new masternodes would not fork even if vote counts were subsequently changed. So while it is technically possible, this is not a trivial exploit and would require a substantial share of the masternodes to achieve. Such an investment would also act as a deterrent to attempt to fork the network.

It is certainly possible to include all information needed for blockchain validation including the votes directly in the blockchain. However, since there is no known attack vector associated with the current system using reasonable assumptions, such a change remains a low priority.

#### 4.2. Multiple Voting

This was reported to Dash and fixed prior to the issuance of the report, and we thank the Veritas Team for finding the bug. It was resolved in PR1135.

#### 4.3. Discarded Votes

If corresponding superblock is in the blockchain already, this wouldn't affect the functioning of the system (see 4.1). If the corresponding superblock is in the future, another finalized budget would be payed in next superblock which is by design. This is also therefore not a confirmed issue.

#### 4.4. Masternode Voting Incentives

True, there is no direct incentive to vote. However as masternode owners are significantly invested in the success of the network, and these operators should care about their investments (i.e., they have indirect incentive to make sure that the right decisions are made by the network), we believe these indirect incentives are sufficient to ensure robust debate and accountability.

We do agree that as the network grows, the current model and granularity of proposals will not scale and will require changes to remain effective. If Dash were to reach Bitcoin's market cap and maintain its existing average proposal size, there would be well over 1,000 proposals monthly. The required time and range of expertise required to evaluate this volume of information would be unrealistic.

The existing system was sufficient to enable rapid growth and fund our most basic needs. However, the team anticipated the scaling issue. It is one of the primary reasons 12.1 is being deployed. 12.1 enables the rapid creation of governance objects that can implement nearly any change the network decides to make.

We will discuss potential approaches to governance from a philosophical and practical standpoint that could provide for the needed scale and expertise to deliver the best outcome for Dash investors. We certainly agree that each masternode operator spending exorbitant amounts of time and independently contracting the needed experts to properly evaluate each proposal would be inefficient and likely lead to ineffective decision-making or non-participation. Indeed, many investors already express voter fatigue.

#### 4.5. Voting Representation

*>As a result, Dash has a system where all decisions are made by a limited set of specific actors. The opinion of the rest of the Dash citizens (miners, regular users) is not considered.*

We don't share the viewpoint that the opinions or needs of miners or regular users are "not considered". Dash's governance system is not and never was intended to be a "representative democracy". We view masternode owners as the owners of the network. Each masternode owner is a significant shareholder, and DGS is a shareholder voting mechanism by design. Masternode owners carrying an owner's mentality are laser focused on the needs of users.

As the "owners" of the system, masternode operators can and should direct the allocation of funds in their own best interests. It was pointed out in section 6.6 that masternode owners' interest in maximizing their

income could lead to “conflicts of interest” such as 1) allocating less of the block reward to miners and more to themselves, 2) increasing transaction fees or block rewards.

However, in the same way that companies can set as high a price as they like on goods or services, market forces will always prevent them from doing so (except at their own peril). Firstly, if Dash’s shareholders spend too little on mining or too little investment in development, and allocate too large a share to themselves, users would leave the insecure and innovation-stagnant product and migrate to better alternatives. Subsequently, the price of Dash would drop, and the value of the masternodes would fall. Secondly, if Dash’s shareholders attempted to overcharge transaction fees, users would again vote with their feet and usage (and total fees collected) would drop, in the same way that a store owner that attempted to price goods at several times the market rate would see sales drop to zero and their profits disappear. Thirdly, if Dash’s shareholders were to increase the block reward, this would erode the trust of users (perhaps permanently), and cause a mass exodus for a more trustworthy alternative.

In fact, Dash “shareholders” are disincentivized from making even responsible changes too rapidly. Consider a scenario in which research determines the needed hashrate expenditures are half of then-current levels. In theory, Dash could reallocate half of the mining reward and fees away from miners instantly without affecting the security of the network, but such a change would demonstrate a disregard for miners’ needs to earn a fair return on their investment, and would stifle future investments in mining. This would obviously not be in the best interest of the network long term. Miners would likely migrate their investments toward networks that showed respect for their business needs, and provided more predictable revenue streams.

In short, free market forces, competition, and common sense constrain the masternodes from behaving irresponsibly. We believe our model incentivizes them to provide greater value to merchants and end-users than other competing networks. In this way, merchant and end-user needs are ensured.

To be clear, we view miners as one of the many “vendors” who are paid to provide the network with a service, albeit important ones. We currently view our users as our “customers”, though we intend to provide an avenue toward ownership in the future (which we will discuss in greater detail below).

#### 4.6. Voting Delegation

Voting Delegation is contemplated in the design of 12.1. However, any changes to DGS (other than technical improvements like combining the multiple superblocs into one superbloc each month, which we did include in 12.1) needs to be approved by the masternodes, rather than implemented directly by the project’s leadership.

#### 4.7. Decreasing Treasury Fund

This criticism is true of the current system. However, at this early stage of development, transaction fees are insignificant, and implementing the 45/45/10 split including the transaction fees requires a fair amount of effort. We anticipate adding the ability to allocate a portion of fees to the treasury at a later date once other higher-priority development efforts are complete.

In the meantime, it only requires 7% growth of the dash/usd rate per year to simply stay at current budget value, and Dash has historically outpaced this rate of growth by a wide margin.

#### 4.8. Flexibility of Treasury Monetary Policy

*>Unspent money from a previous cycle can’t be transferred to the next cycle, and any funds unspent are simply never created (which has a similar effect to burning funds).*



This has been contemplated and in fact we have stated that a future protocol change could even permit the unspent funds from previous cycles to be recaptured. This would require the approval of the network and a fair amount of effort, so it is further down the list of priorities to address.

*Evan: I like the idea of using burn funds in this process as a network savings account. The current implementation could be easily altered to bring back the burnt coins, in a one time transaction for the future, using a special type of superblock. This means we could volunteer to burn coins, when the projects don't look great on the docket and "save" for a rainy day... It doesn't affect the long term emission rate either.*

In principle, we agree that more flexibility would be useful. 10% is not a magic number, and presumably if masternode owners believe that faster development could result from some incremental allocation beyond 10% they believed would yield better overall returns, that should be made possible.

At some point in the future, scientific research into the allocation to mining required to secure the network may permit the network to safely reallocate a portion of the mining reward elsewhere. As incremental mining hashrate offers diminishing returns after extremely high security is already achieved, it is logical that any "waste" above the level required would be better allocated to more productive activities or returned to the masternode owners as higher return if no productive use is found. The masternode owners could even determine that reducing the block reward, rather than reallocating it would produce better results by effectively reducing the dilutive "tax" inflation represents to users.

Ultimately, these types of changes can yield a more cost efficient, lower "tax", and/or more effective network.

>What if the Dash community grows and more funding is needed beyond what the current 10% can provide?

If the Dash community grows, the value of each coin would likely grow in direct proportion. We anticipate any network funding needs are either fixed costs, or at worst grow linearly with demand, so it's unlikely that network growth would be the culprit for future funding shortfalls. If anything, network growth should provide budget pressure relief.

#### 4.9. Code Review

We thank the authors of the report for their review. The code is being actively developed so expect more bugs to be identified, especially for any brand new code. Message propagation bugs have constituted the most time consuming to resolve, but over 85% of the root causes have now been identified and resolved as of this writing. As always, we welcome everyone to help us to find and squash bugs and make our code base as stable and robust as possible before release.

## Proposed Solutions

In this section we will address the merit of various proposed solutions contemplated by the Veritas Team not otherwise addressed above.

### 6.2.1. Flexible Participation Requirements

We agree, and an even more flexible one than the one proposed is already on the roadmap. Evolution "savings accounts" are effectively decentralized masternode shares, which will support much smaller investments that earn a return. We have publicly referred to these participants as "owners of a mutual bank". Mutual banks are financial entities in which the customers are also the owners.



### 6.3. Implementation of Delegation with Liquid Democracy

We agree, but this is far from the only model which can address scaling and expertise requirements while maintaining decentralized control and authority over the system. The possibilities are too wide-ranging to list, but differing models will require evaluating the inherent trade-offs involved. It is not clear that the proposed solution is the best of all alternatives.

### 6.4. Assigning Categories to Each Funding Proposal

We strongly disagree with this model on the basis of practical experience. On the surface, this sounds like a logical and effective solution. However, as the authors point out “there are no universal experts” and unfortunately, proposals are extremely unlikely to fall neatly into anything resembling the example provided, especially as the network grows. Additionally, effective decision-making requires coordination across the “silos” proposed.

Consider that all large financial organizations are matrixed, meaning they have at least two of the following dimensions embedded in the organization: geography, function, product, market, or distribution channel. The exact makeup of the structure is informed by the needs of the organization, which can change over time. For example, a global bank might organize itself primarily by geography, if the competitive landscape dictates different approaches in each region. They might then have a product group that works across each region to customize products to meet each region’s needs. Structures like these exist primarily because delivering the most effective solutions requires maximizing expertise (knowledge and skill) in all geographies and products simultaneously, which is unlikely to be most effectively developed in one individual or group.

Now consider a proposal far off in the future to create a new product for the European market through Dash’s “virtual teller machines” distribution channel. Which committee should evaluate this proposal? Should it go to the infrastructure committee that oversees the virtual teller machines, or will they not understand the product or European market well enough? These types of decisions require human judgement to bring together a combination of the right expertise and resources (e.g., research, vendors, etc) from various areas based on each situation. We believe it is naive to assume static groups will be able to best address each opportunity. Category assignment definitely simplifies the selection of experts (as pointed out), but likely comes at the direct expense of effectiveness.

Consider also the effectiveness of decision-making in silos. Most decisions require a strategy and coordination across functions. For example, imagine the development team approves a grand new service, yet the infrastructure team decides against expansion of its web server capacity needed to handle the change, and the marketing team decides to double-down on marketing existing services. This is obviously far from ideal and separate committees making decisions independently doesn’t facilitate the coordination required to compete effectively long-term. Networks that embrace a fragmented decision-making framework are likely to be outperformed over time by networks that embrace a coordinated strategy framework.

### 6.5 Introduction of Paid Expert Roles

A person who is paid to provide their judgement and expertise and make recommendations and act in the best interest of the network sounds an awful lot like an employee to us. There doesn’t seem to be any improvement in decision-making outcomes in the model proposed over a more traditional “employee” model. However, we do see complications:

1. The “experts” making the decisions do not appear to be the employees responsible for implementation. This creates a hazard... if the project fails, who is to blame? The “expert” would naturally point to the employees and claim the idea was implemented poorly. The “employees” or contractors would point to the expert and claim the idea was poorly conceived, no matter how well it was implemented, and never should have been approved. Ultimately, neither may be held accountable. Within our own structure / DAO, we prefer a model in which the decision-making and implementation are conducted by the same entity (e.g., “one throat to choke”), isolating responsibility for successes and failures.
2. Additional overhead is introduced. Presumably, the individuals working for the network should be experts in their field already - be it programming, product development, etc. This calls into question the need to introduce other experts to make decisions on behalf of the network that these individuals are presumably capable of already.

For these reasons, we feel the proposed solution is ill-conceived and requires a better explanation for the potential benefits of this additional layer of experts.

That said, we **do** see a need for panels of experts for very deep research into specific topics that exceed the scope or capabilities of the organization(s) and individuals supporting the network. This may be because the need for these experts is not permanent (one-time questions), addresses a cross-disciplinary question (brings together diverse backgrounds for a stated purpose), requires concentrated/prolonged research in new fields (and don’t want to take experts away from existing duties), or requires independence from the organization itself (e.g., recommendations about compensation changes).

## 6.6 Expanding List of Voting Entity Types

We do not share the view that voting rights are best allocated to all involved parties. We believe the masternode operators - with an ownership mindset - are most effectively positioned to make decisions in the best interest of a healthy system. The Veritas Team points to the potential “conflicts of interest” of these owners as a flaw, envisioning “a usurpation of cryptocurrency control. This, in turn, can lead to a loss of other participants’ interests, leading them to leave the system and destroying the cryptocurrency itself.” That simply hasn’t happened yet, and for the free-market reasons described in section 4.5 above, we believe it is extremely unlikely heavily invested masternode owners would behave recklessly enough to destroy their own investment.

At the same time, the analysis neglected to point out the potential perils of end-user voting or miner voting. We view these conflicts of interest as very real, unlike masternode “conflicts”. Consider end-users voting on fee reductions. They would always be incentivized to vote for fee reductions at the expense of masternode owners no matter how low rates became. They might also vote only to spend funds on services they foresee using, even if a subset of other users would value them deeply. There are no natural market forces that compel them to do otherwise because they are not “owners”. Likewise, miners would always vote for proposals that allocated more to miners at the expense of masternodes. Again, there are no natural market forces that compel them to act differently.

So we fundamentally disagree with allowing “customers” and “vendors” (not just miners, but other contractors as well) to make decisions about our “business”.

All that said, we do plan to implement a form of end-user voting eventually through the Evolution savings accounts (e.g., masternode shares). These end-user funds are not “transitory” in nature... they are by definition long-term holdings of Dash, which we feel qualifies them as “owners” of our “mutual network”. To the extent that the user is invested, they are more heavily incentivized to make decisions in the best interest of the network itself, rather than their own self-interest.

If any other group is worthy of inclusion, perhaps it should be the “employees” - the men and women that work to improve the network day after day. Many businesses seek input from and consider the impact on their employees when making decisions, and we would applaud an approach to incorporate that thinking into a governance model, which we believe could have merit by helping to create a better working environment, which could have a broad range of impacts from employee retention to end-user satisfaction.

## Other Weaknesses

Practical experience with our governance system has revealed other weaknesses as well which we plan to resolve. We list them here with little comment, but wanted to capture them.

1. The inability to cancel a proposal that is no longer desired. Today, proposal owners must request a “downvote” if a proposal is no longer desired, which cannot be achieved quickly or without risk of undesired consequences.
2. The inability to modify a proposal (e.g., reduce the amount requested) without resubmitting.
3. The inability to fix the fiat equivalent of multi-month payments (which might be desired for fiat subscriptions for example)
4. The inability to enter into “contracts” that are prioritized to ensure fulfillment of all future payments once approved. An example is a three month contract that although it has sufficient support in month 1 and 2, fails to fund in month 3 because of more competition that month (but would otherwise fund in a less-competitive month), resulting in a breach of contract.

## Summary

The Veritas Team put forward a number of potential improvements to the DGS, which helped stimulate discussion within the team and quite frankly forced us to articulate in greater detail our vision and approach to governance. Where the team already had plans in the roadmap for similar improvements, it was validating to see similar ideas developed independently from our own. Where we disagree with proposed improvements, we frequently viewed these differences as “design choices” rather than “flaws”, and we recognize other projects will likely favor other design choices. Ultimately, competition will determine which model(s) yield superior results and attract users and investors.

However, we also want to make clear that these views are the views of the core team, and one of our deeply-held values is that any of our proposed “design choices” we implement to governance in the future should be validated by the network itself. The network may trust our judgment, value our track record, and implement changes we recommend. Or it may decide to go a different route by rejecting some or all of the proposed changes. Either way, we believe it is the masternodes’ decision to make, and any change should remain revocable by the network.



***Dash is Digital Cash***

[www.dash.org](http://www.dash.org)