



## Improvising the Global Naming System using Blockchains

Jasmine Bedi Khurana, Vivek Tomar, Utkarsh Wadhwa

Information Technology, Galgotias College of Engineering & Technology,  
Uttar Pradesh, India

**Abstract**— *In the past few years, cryptocurrency networks like Namecoin, Bitcoin and other peer to peer networks have shown confidence in them by using naming systems with no trusted parties operating them. Different users can associate them with the system using human readable names and only the true owner of the private key can modify the information. This paper is more concerned about our experience of the real-world implementation of the PKI service on Namecoin network. We want to share our concern on the various problems like reliability, security issues, etc when we are registering more than 33k entries which involved about 200k transactions on the network of Namecoin. There is a new naming system named as Blockstack which is an open source system and involves about 40k users under PKI system. We later migrated to the Bitcoin system from Namecoin and the experience is shared in this paper.*

**Keywords**— *Cryptocurrency; Digital Currency; Bitcoin; Blockchain; Cryptography; 51% attack; Blockstack; DDOS attack.*

### I. INTRODUCTION

Cryptocurrency blockchains and their respective P2P networks are useful beyond exchanging money. They provide cryptographically auditable, append-only ledgers that are already being used to build new, decentralized versions of DNS and public-key infrastructure (PKI), along with other applications like file storage [11] and document timestamping [10]. Because blockchains have no central points of trust or failure, they enable a new class of decentralized applications and services that minimize the degree to which users need to put complete trust in a single party, like a DNS root server or a root certificate authority.

Blockchain networks have attracted a lot of interest from enthusiasts, engineers, and investors; 955 million USD has been invested in blockchain start-ups in the last 3 years [9]. With the rapid capital infusion, infrastructure for blockchains is getting quickly deployed [8] and blockchains are emerging as publicly-available common infrastructure for building decentralized systems and applications. However, blockchain networks are at a very early stage and there is very little production data available to guide design trade-offs.

Many non-financial applications of blockchains imply the need for a naming system that securely binds names, which can be human-readable, to arbitrary values. The blockchain gives consensus on the global state of the naming system and provides an append-only global log for state changes. Writes to name/value pairs can only be announced in new blocks, as appends to the global log. The global log is logically centralized (all nodes on the network see the same state), but organizationally decentralized (no central party controls the log).

The decentralized nature of blockchain-based naming introduces meaningful security benefits, but certain aspects of contemporary blockchains present technical limitations. Individual blockchain records are typically on the order of kilobytes and cannot hold much data. Latency of creating/updating records is capped by the blockchain's write propagation and leader election protocol, and it is typically on the order of 10-40 minutes [7]. The total new operations in each round are limited by average bandwidth of nodes participating in leader election (for Bitcoin the current bandwidth per new round is 1MB). Further, new nodes need to independently audit the global log from the beginning, which implies that as the system makes forward progress, the time to bootstrap new nodes increases linearly.

We believe that in spite of these scalability and performance challenges, blockchains provide important infrastructure for building secure, decentralized services. The cost of tampering with blockchains grows with their adoption; today, it would require hundreds of millions of dollars to attack a large blockchain like Bitcoin [1].

These benefits motivated us to use blockchains to build a new decentralized PKI system. Our system enables users to register unique, human-readable usernames and associate public-keys, like PGP, along with additional data to these usernames. There is no need for any central or trusted party in our PKI system. This paper presents our experiences from operating this PKI system on the Namecoin network, which is one of the largest services built on top of a blockchain to date. We outline the challenges that we had to overcome for registering/updating over 33,000 user entries and for sending over 200,000 transactions on the Namecoin network.

Our production deployment led to many interesting experiences where we observed and analysed network anomalies and security problems that were not discovered or documented before. We discovered a critical security problem where a single miner consistently had more than 51% relative compute power, called mining power, on Namecoin (see [18] for details on the 51% attack and mining power). A 51% attack is one of the most serious attacks on a blockchain and impacts the security and decentralization properties.

Moreover, we also encountered chronic networking issues with broadcasting transactions on the Namecoin network. Reliability of the network generally depends on how actively a blockchain network is monitored and maintained, as well as financial incentive for operating the network. Therefore, for both security and reliability reasons, blockchain-based services should use the largest and most-secure blockchain, which is currently Bitcoin.

Our experience with Namecoin informed the design and implementation of a new blockchain-based naming system, called Blockstack that uses the Bitcoin blockchain. Unlike previous blockchain-based systems, Blockstack separates its control- and data-plane considerations: it keeps only minimal metadata (namely, data hashes and state transitions) in the blockchain and uses external data stores for actual bulk storage. Blockstack enables fast bootstrapping of new nodes by using check-pointing and skip lists to limit the set of blocks that a new node must audit to get started. We have released Block-stack as open source [6].

## **II. MOTIVATION AND BACKGROUND**

In this section, we describe the motivation for building naming systems that have no central point of trust and provide the relevant background on blockchains. In this paper, we use the term naming system to mean a) names are human-readable and can be picked by humans, b) name/value pairs have strong sense of ownership—that is, they can be owned by cryptographic keypairs, and c) there is no central trusted party or point of failure. Building a naming system with these three properties was considered impossible according to Zooko's Triangle [17] and most naming systems provide two out of these three properties [16]. Namecoin [37] used a blockchain-based approach to provide the first naming system that offered all three properties: human-readability, strong ownership, and decentralization.

### **A. Background on Blockchains**

Blockchains provide a global append-only log that is publicly writeable. Writes to the global log, called transactions, are organized as blocks and each block packages multiple transactions into a single atomic write. Writing to the global log requires a payment in the form of a transaction fee. Nodes participating in a blockchain network follow a leader-election protocol for deciding which node gets to write the next block and collect the respective transaction fees. Not all nodes in the network participate in leader election. Nodes actively competing to become the leader of the next round are called miners. At the start of each round, all miners start working on a new computation problem, derived from the last block, and the miner that is the first to solve the problem gets to write the next block. In Bitcoin, the difficulty of these computation problems is automatically adjusted by the protocol to get 1 new block roughly every 10 minutes. See [7] for further details on how blockchains work.

### **B. Namecoin's Naming System**

Namecoin is one of the first forks of Bitcoin and is the oldest blockchain other than Bitcoin that is still operational, with a current market cap of 5 million USD [4] (the market cap of a cryptocurrency is the exchange-traded value of its coins multiplied by its number of coins in existence). The main motivation for starting Name-coin was to create an alternate DNS-like system that re-places DNS root servers with a blockchain for storing information on registered domain names [37]. Given that blockchains don't have central points of trust, a blockchain-based DNS is much harder to censor and registered names cannot be seized from owners without get-ting access to their respective private keys [16] Altering name registrations stored in a blockchain requires prohibitively high computing resources because re-writing blockchain data requires proof-of-work [5].

Namecoin is derived from Bitcoin's code, and keeps most functionality identical to Bitcoin, with the exception of support for registering name/value pairs. Our work on Blockstack implements name registration functionality as a separate layer on top of Bitcoin, but this design was non-obvious before our work. Indeed, it was common practice to start alternate blockchains by forking them from Bitcoin and make modifications required by the respective service/application, the precise approach taken by Namecoin.

Just like DNS, there is a cost associated with registering a new name. The name registration fee discourages people from grabbing a lot of names that they don't in-tend to actually use. In Namecoin, the recipient of registration fees is a "black hole" cryptocurrency address from which money cannot be retrieved [16]. Namecoin defines a pricing function for how cost of name registrations change over time. Namecoin supports multiple namespaces (like TLDs in DNS), and the same rules for pricing and name expiration apply to all namespaces. By convention, the d= namespace is used for domain names. For example, to register the domain yahoo on Namecoin, one must register the name d=yahoo and then put the IP address of the Yahoo! website in the name/value pair.

In Namecoin, name registration is a two-step process. A user first pre-orders a name in a new transaction that includes hash (name) in the transaction. This does not reveal what name she is trying to register. After the pre-order transaction, has been confirmed by the network— i.e., enough blocks (typically 10) are later added to the blockchain to make it computationally infeasible for any miner to re-write recent blockchain history and reverse the transaction—the user can reveal the name she was actually trying to register. This is done by sending a second transaction on the network that completes the register step. The user includes the value of the name/value pair in the second transaction as well. The cryptocurrency address that signed the two transactions becomes the owner of the newly registered name/value pair.

Name registrations expire after a fixed amount of time, measured in new blocks written (currently 36,000 blocks, which translates to roughly 8 months). Namecoin also supports updating the value associated with a name, as well as ownership transfers.

**C. Blockchain-based PKI system**

We used Namecoin to build a PKI system, called Blockchain IDs, by starting a new namespace u= on it. We defined the format for publishing public-keys, like PGP, along with other profile data in the blockchain [2]. This is similar to the format of DNS records. We launched a web service in March 2014 that enabled people to easily register names on the u= namespace of Namecoin and associate profile data with them. In our web service, we first register the name on the user’s behalf (and also cover the cost of name registration for them) and then transfer the name to a cryptocurrency address owned by the user. Our implementation is the first PKI system that binds user identities to public-keys using a blockchain. All registered names have an ECDSA public-key [13] binding by default, and a subset of users added their PGP keys as well. More than 33,000 users registered on the Namecoin blockchain between March 2014 and August 2015 through our service. According to a study by Harry et al. [16], our system is the second largest namespace on Namecoin by volume and the largest by number of active users.

**III. NAMECOIN DEPLOYMENT EXPERIENCE**

The security of name ownership is tied to the security of both the underlying blockchain and the software powering it. There are three security issues to consider:

**A. Security**

1) Cost of Attack: Miners often pool their resources to form a mining pool, which is essentially a super node on the network (a lot of computational power behind a single miner node). If the amount of computational power under the control of a single miner (or pool) is more than the rest of the network, called a 51% attack, then that miner has the ability to attack the network and rewrite recent blockchain history, censor transactions (e.g., for name registrations), and steal cryptocurrency using double spend attacks. This is because it will win the leader election majority of the time, and produce a blockchain history with more proof-of-work than any disagreeing miner. The more expensive it is to control a majority of the compute power on a particular blockchain, the more secure the blockchain.

We noticed in late 2014 that a single mining pool consistently had more than 51% of the compute power on Namecoin. Recently, the situation has been even worse, with a single mining pool controlling over 60% of Name-coin’s compute power. Figure 1 shows the monthly, weekly, and daily distribution of mining power for the month of August 2015, right before we migrated our system away from Namecoin. In fact, we have observed F2Pool (also known as Discus Fish) control up to 75% compute power in a particular week. At such concentration, Namecoin is effectively controlled by a single party; F2Pool gets to write most of the new blocks and can undermine the security of the blockchain at will.

2) Software Vulnerabilities: Raw hashing power is not the only metric for the security of a blockchain. Software issues/bugs are also very important, e.g., a Namecoin bug allowed people to steal names from any-one [12]. Actively developed codebases with frequent security reviews lower the rate of critical bugs in production. Despite originally sharing the same codebase as Bitcoin, Namecoin was not kept up-to-date with advancements made in Bitcoin development.

3) DDoS Attacks: One of the less explored security issue with cryptocurrency blockchains is network at-tacks, e.g., DDoS attack on core discovery nodes, mining pools, or the entire network. The more peers a cryptocurrency network has, the more resilient the network is to denial-of-service attacks. As a relatively small net-work, Namecoin has many fewer nodes than Bitcoin (170 vs. 4,600 in Jan 2016 [3]), which makes it more vulnerable in this respect.

When considering security on these three metrics, the Bitcoin blockchain is currently by far the most se-cure blockchain. We faced a fundamental trade-off between security and introducing new functionality to blockchains. Starting

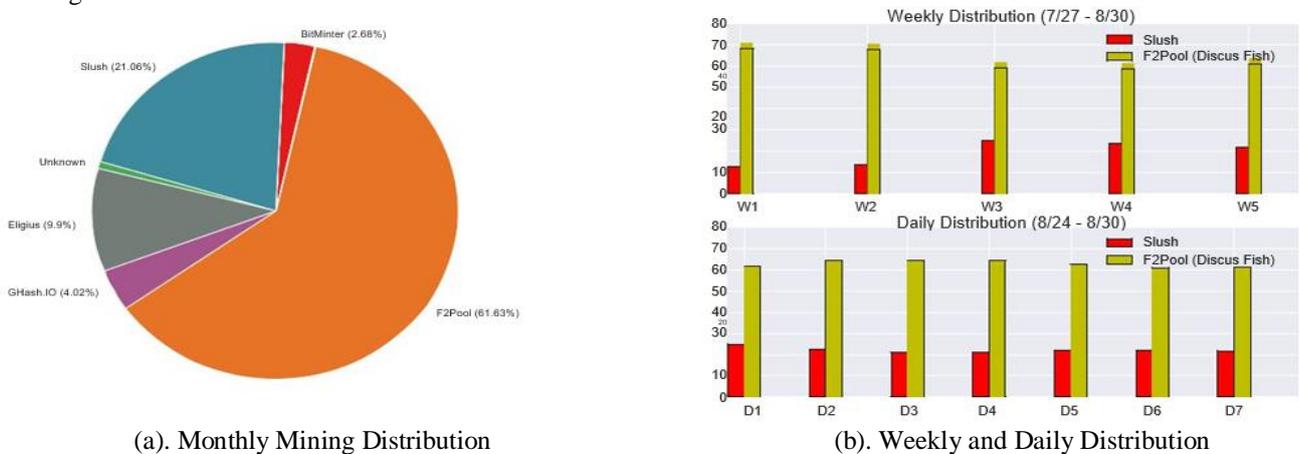


Figure 1 (Mining Distribution)

A new blockchain network is how developers introduce new functionality not provided by Bitcoin e.g., a naming system that is of interest to many emerging applications. However, new blockchains are significantly less secure than Bitcoin. We introduce Blockstack to show how to overcome this trade-off by creating virtual chains to introduce new functionality as a layer on top of Bitcoin.

## B. Network Reliability and Throughput

The throughput of our PKI system (number of entries we can register/update) is directly dependent on the through-put of the underlying blockchain. The number of new register/update operations that can be performed per hour is limited by the number of transactions that can be sent (and confirmed) on the underlying blockchain per hour. Similarly, reliability of our PKI system is impacted if the underlying blockchain cannot perform operations reliably and consistently.

**Network Latency Spike:** As a fork of Bitcoin, Name-coin shares many protocol properties with Bitcoin, including a 10 minute target of writing new blocks (latency target) and a 1MB bandwidth limit on block size (giving throughput of 1000 transactions per block). Figure 2(a) shows that since we launched our PKI system in March 2014, Namecoin on average performed well on the net-work latency target. As expected, most new blocks were written within 10 and 40 minutes (similar times have also been observed on Bitcoin [7]). Figure 2(b) shows an incident in late August 2014 (at block number 192000), where network latency skyrocketed for a couple of weeks (1000 blocks are roughly a week). After investigating the issue and having discussions with Namecoin developers, we discovered that the latency spike was caused by software issues in Namecoin. Someone on the net-work was sending transactions with a large number of data fields per transaction. This was causing severe performance problems for the miners and their Namecoin daemons kept crashing. Without stable miner nodes there is no one there to package blocks, and hence the longer delay in new blocks. This shows that unexpected protocol/software issues can trigger network latency problems. During this period, we noticed a slowdown in rate of new registrations of our PKI system along with a spike in user complaints and support requests.

**Network Throughput Drop:** In early September 2014, right after the latency spike incident we noticed that our transactions were not getting accepted for many consecutive blocks and, after a while, will get accepted in bulk in a single block that packaged a lot of transaction. We noticed that a lot of new blocks had no trans-actions in them. This issue persisted for over a week and Figure 3 plots the number of transactions that we were trying to send (shown as “tx target”) vs. the number of transactions that were getting accepted by the network. Network latency was completely normal (shown at top of Figure 3), but network throughput went down because of no transactions in new blocks. We tried upgrading our software and rebroadcasting transactions, but the issue persisted. We concluded that there is a large mining pool that is either intentionally refusing or is unable to pack-age transactions in the blocks it is writing. Our transactions will get packaged only when some other miner was elected to write the new block.

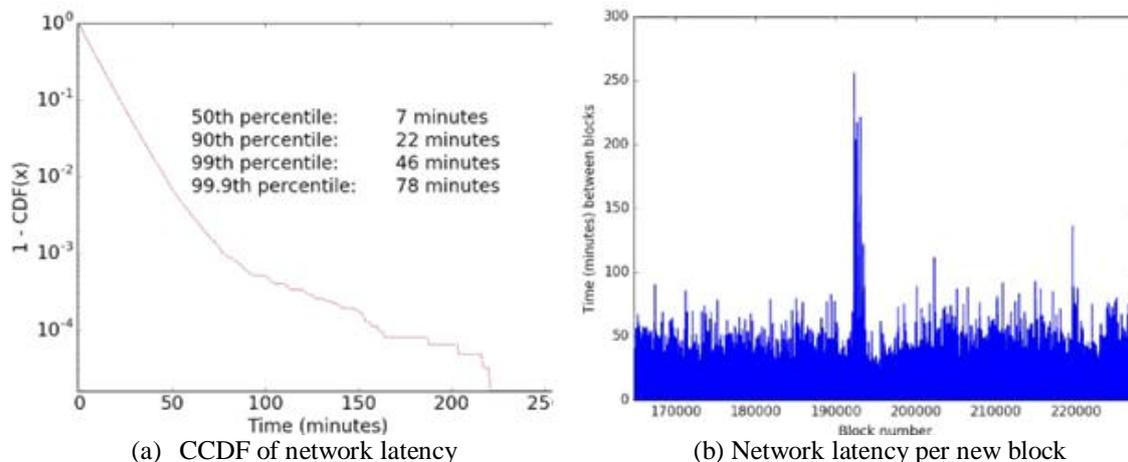


Figure 2 (Network Latency)

## C. Failure of Merged Mining

Security of a blockchain depends on relative compute power that miners have and how much would it cost a single party to get more computing power than the rest of the network. New, smaller blockchains have a boot-strapping problem where in the initial days of a new blockchain it would be relatively easy for a single party to take it over, since the total compute power on the blockchain is not yet large enough to prevent this.

To address this problem Satoshi Nakamoto, author of Bit-coin, introduced “merged mining” where an alternate blockchain can allow Bitcoin miners to participate in the new network without requiring them to spend extra compute cycles. The miners can make extra profits on the new blockchain without adding computational overhead. With a merge-mined cryptocurrency the security of the blockchain is typically a subset of the “main blockchain”, because in practice not all miners of the main blockchain go through the trouble of setting up merged mining.

Namecoin switched to merged mining with Bitcoin to increase security of the blockchain [16]. Name-coin is the oldest and largest merged-mined cryptocurrency and inspired other cryptocurrencies to consider it as well. One of our key findings is that merged mining is currently failing in practice: the leading merged-mined blockchain, Namecoin, is vulnerable to the 51% attack. Moreover, merged-mining provided a false sense of security to the Namecoin community. If Namecoin shares compute power of Bitcoin that might be a high compute power in absolute terms but a single miner can still have more than 51% compute power on the merged-mined blockchain. The security of the blockchain depends on relative compute power and not on absolute compute power. F2pool, for example, had 30-35% compute power of Bitcoin, when it had more than 60% compute power of Namecoin. Unless the merged mined cryptocurrency can

consistently get a very high ratio of main blockchain miners to support their software, merged mining will not keep it safe from 51% attacks.

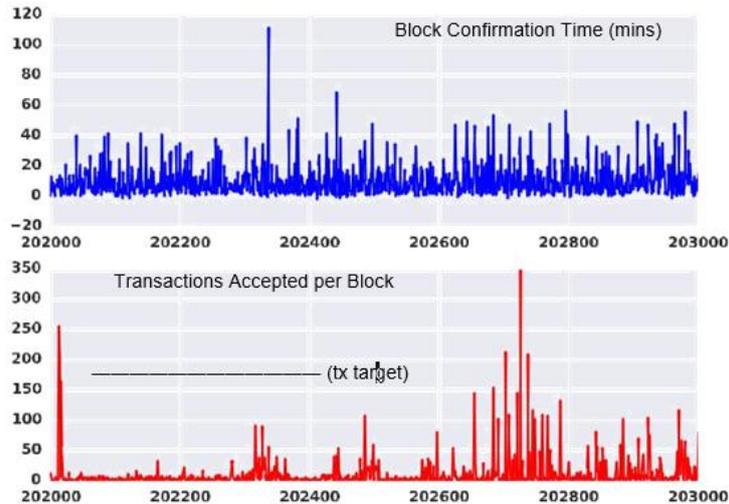


Figure 3 (Throughput, 'tx target' is the number of trans-actions we were trying to send)

The merged mining failure of Namecoin convinced us that we're at a stage in the evolution of blockchains where there is not yet enough compute cycles dedicated to mining to support multiple secure blockchains. The respective financial capital attached to blockchains relative to Bitcoin supports this argument: Bitcoin has a 5.9 billion USD market cap, which ac-counts for 89% of the market cap of all 500+ blockchains combined, while the 2nd and 3rd largest market caps are 3.2% and 2.6% of Bitcoin respectively [5]. It's possible that after some years or decades we might have multi-plea secure blockchains, but in the near-future Bitcoin's blockchain is the only one that is prohibitively expensive to attack.

#### IV. DESIGN OF BLOCKSTACK

Blockstack is designed to implement a naming system with human-readable names in a layer above the blockchain. In this section, we describe how Blockstack uses the underlying blockchain, and present how it copes with technical limitations of contemporary blockchains.

##### A. Challenges

Building systems with blockchains presents challenges:

**Limits on Data Storage:** Individual blockchain records are typically on the order of kilobytes and cannot hold much data. Moreover, the blockchain's log structure implies that all state changes are recorded in the blockchain. All nodes participating in the network need to maintain a full copy of the blockchain, limiting the total size of blockchains to what current commodity hard-ware can support. As of January 2016, Bitcoin nodes need to dedicate 53GB total disk space to blockchain data for staying synchronized with the network.

**Slow Writes:** The transaction processing rate is capped by the blockchain's write propagation and leader-election protocol, and it is pegged to the rate at which new blocks are announced by leader nodes (called miners in blockchain networks [7]). New transactions can take several minutes to a few hours to be accepted.

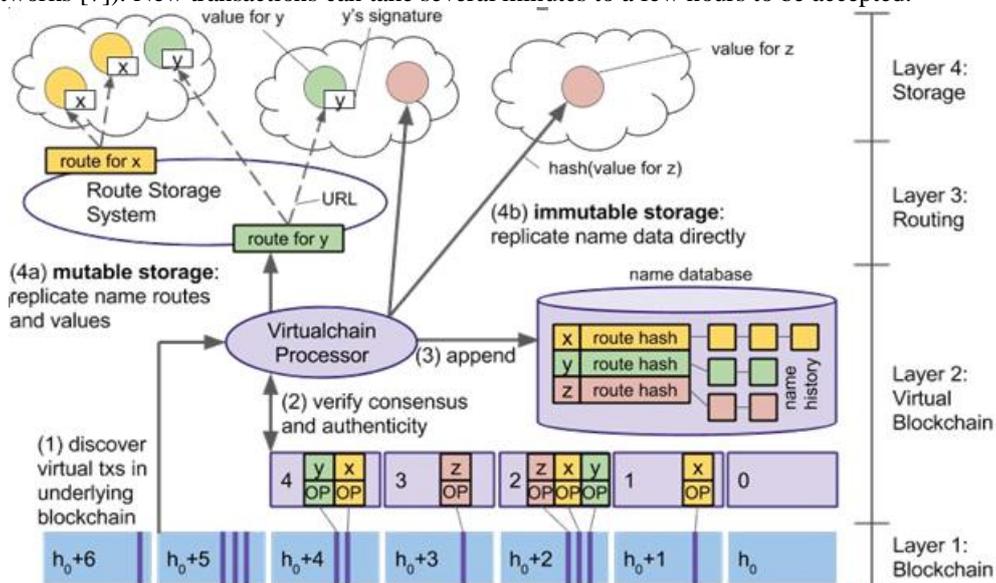


Figure 4 (Overview of Blockstack's architecture)

**Limited Bandwidth:** The total number of transactions per block is limited by the block size of blockchains. To maintain fairness and to give all nodes a chance to become leader in the next round, it's required that all nodes receive a newly announced block at roughly the same time. Therefore, the block size is typically limited by average uplink bandwidth of nodes in the network [7]. For Bitcoin the current bandwidth is 1MB (1000 transactions) per new block.

**Endless Ledger:** Integrity of blockchains depends on the ability for anyone to audit them back to the first block. As the system makes forward progress and issues new blocks, the cost of this auditing grows linearly with time and booting up new nodes becomes progressively more time consuming. We call this the end-less ledger problem. Bitcoin's blockchain currently has 395,000 blocks and new nodes take 1-3 days to download blockchain from Bitcoin peers, verify it, and bootstrap.

## **B. Architecture Overview**

Blockstack builds a naming system as a separate logical layer on top of the underlying blockchain. Blockstack uses the underlying blockchain to achieve consensus on the state of this naming system. It uses the underlying blockchain as a communication channel for announcing state changes; any changes to the state of the naming system can only be announced in new blockchain blocks. Relying on the consensus protocol of the underlying blockchain, Blockstack can give total ordering for all operations (like register, update, and transfer) supported by the naming system.

**Separation of Control and Data Plane:** Blockstack decouples security of name registration and name ownership from data availability of values associated with names. This is done by separating the control and data planes (Figure 4). The control plane is responsible for registering human-readable names and creating (name; hash) bindings. It also defines the protocol for establishing ownership of names, which are owned by cryptographic keypairs. The control plane consists of a cryptocurrency blockchain and a logically separate layer on top, called a "virtual blockchain".

The data plane is responsible for data storage and availability. It consists of (a) routes for discovering data, and (b) external storage systems for storing data (such as Amazon S3, IPFS [14], or Syndicate [15]). Data values are signed by public keys of respective owners. The control plane retrieves data values from the data plane and verifies their authenticity by checking either the hash of the data or the signature of the public key.

We believe this separation is a significant improvement over Namecoin, which implements both control and data plane at the blockchain level. Our design not only significantly increases the data storage capacity of the naming system, but also allows each layer to evolve and innovate independently of each other.

**Agnostic of Underlying Blockchain:** The design of Blockstack does not put any limitations on which cryptocurrency blockchain can be used with it. Any blockchain can be used, but the security and reliability properties are directly dependent on the underlying blockchain. We believe that the ability to migrate from one blockchain to another is an important design choice as it allows for the larger system to survive, even when the underlying blockchain is compromised. Currently, Blockstack core developers decide which underlying blockchain(s) to support in which version of the software. Individual applications can decide to run the software version of their choice and keep their namespace on a particular blockchain, if they prefer to not migrate.

**Ability to Construct State Machines:** A key contribution of Blockstack is the introduction of a logically separate layer on top of a cryptocurrency blockchain that can construct an arbitrary state machine after processing information from the underlying blockchain. We call this layer a virtual blockchain. A virtual blockchain treats transactions from the underlying blockchain as inputs to the state machine. Valid inputs trigger state changes. At any given time, where time is defined by the block number, the state machine can give exactly one global state. Time moves forward as new blocks are written in the underlying blockchain and the global state is updated accordingly.

Virtual blockchain can introduce new types of state machines without requiring any changes from the underlying blockchain. Introducing new state machines directly in a blockchain requires, potentially consensus breaking, upgrades of the cryptocurrency blockchain that can cause forks and are hard to do in practice [7]. Currently, Blockstack introduces a state machine that represents the global state of a naming system i.e., who owns a particular name, what data is associated with a name etc. It's possible to use the virtual blockchain concept to define other types of state machines as well.

## **C. Blockstack Layers**

Blockstack introduces new functionality on top of blockchains by defining a set of new operations that are otherwise not supported by the blockchain. Blockstack has four layers, with two layers (cryptocurrency blockchain and virtual blockchain) in the control plane and two layers (routing and storage) in the data plane.

### *Layer 1: Cryptocurrency Blockchain*

The blockchain occupies the lowest tier, and serves two purposes: it stores the sequence of Blockstack operations, and provides consensus on the order in which they were written. Blockstack operations are encoded in transactions on the underlying blockchain.

### *Layer 2: Virtual Blockchain*

Above the blockchain is a virtual blockchain, also called virtual chain. The main advantage of virtual chains is to introduce new functionality/operations without requiring changes to the underlying blockchain. Only Blockstack nodes are aware of this layer and underlying blockchain nodes are agnostic to it. Blockstack operations are defined in the virtual chain layer and are encoded in valid blockchain transactions as additional metadata. Blockchain nodes do see the raw transactions, but the logic to process Blockstack operations only exists at the virtual chain level.

The rules for accepting or rejecting Blockstack operations are also defined in the virtual chain. Accepted operations are processed by the virtual chain to construct a database that stores information on global state of the system along with state changes at any given blockchain block. Virtual chains can be used to build a variety of state machines. Currently, Blockstack defines only a single state machine; a global naming system.

*Layer 3: Routing Layer*

Blockstack separates the task of routing requests (i.e., how to discover data) from the actual storage of data. This avoids the need for the system to adopt a single or particular storage service from the onset, and instead allows multiple storage providers to coexist and compete for use (including both commercial services and peer-peer systems).

The virtual chain binds names with respective hash (route) and stores these bindings in the control plane, whereas the actual routes are stored in the routing layer. Users do not need to trust the routing layer, because the integrity of routes can be verified by checking the hash (route) in the control plane.

In Blockstack’s current implementation, nodes form a DHT-based peer network for storing routes. The DHT only stores routes if hash (route) was previously announced in the blockchain and effectively white-lists the data that can be stored in the DHT. Due to space constraints, we omit most details of our DHT storage from this paper; the key aspect relevant to the design of Blockstack is that routes (no matter from where they are fetched) can be verified and therefore cannot be tampered with.

*Layer 4: Storage Layer*

The top-most layer is the storage layer, which hosts the actual data values of name/value pairs. All stored data values are signed by the key of the respective owner of a name. By storing data values outside of the blockchain, Blockstack allows values of arbitrary size, and allows for a variety of storage backend. Users do not need to trust the storage layer, because they can verify the integrity of the data values in the control plane. There are two modes of using the storage layer and they differ in how the integrity of data values is verified; Blockstack supports both storage modes simultaneously.

a) Mutable Storage is the default mode of operation for the storage layer. Bindings between name and hash (route) are kept in the control plane and the routing layer is used to discover data values. Data integrity is verified by validating the signatures associated with values. Mutable storage allows for faster writes, since data updates do not involve any transactions on the blockchain, which are slow. Updates to name/value pairs do not involve the blockchain.

b) Immutable Storage by-passes the routing layer and stores bindings between name and hash (data) in the control plane, instead of bindings between name and hash (route). Data integrity is verified by hashing the data and comparing the result against the hash (data) from the control plane. This mode is suitable for data values that don’t change often and where it is important to verify that you are viewing the latest version of the data value. For immutable storage, updates to data values require a new transaction on the underlying blockchain, making data updates much slower than mutable storage.

**D. Naming System**

Blockstack uses the four-tiered system to implement a complete naming system. Names are owned by cryptocurrency addresses of the underlying blockchain and the associated private keys e.g., ECDSA-based private keys used in Bitcoin [7]. As with Namecoin, a user pre-orders and then registers a name in two steps, in order to claim a name without revealing it to the world first (otherwise an attacker can race the user in claiming the name). The first user to successfully write both a pre-order and a register transaction owns the name. Any previous preorders become invalid when a name is registered. Once registered, a user can update the name/value pair by sending an update transaction (which changes the name/value binding) and uploading the new value to the storage layer. Name transfer simply changes the cryptocurrency address that is allowed to sign subsequent transactions revoke disables any further operations.

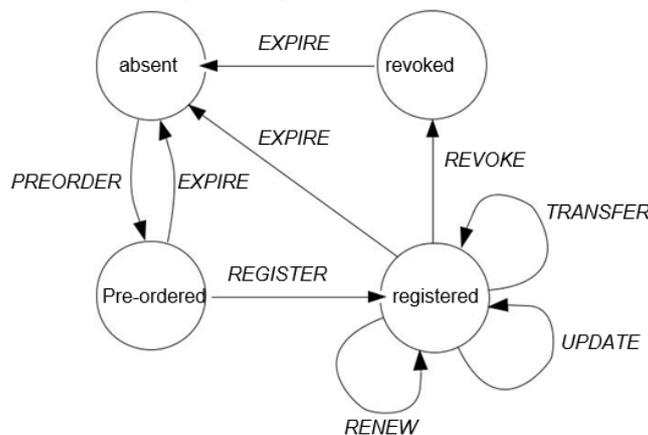


Figure 5 (States and Transitions for a Name)

The naming system is implemented by defining a state machine and rules for state transitions in the virtual chain. Figure 5 shows the different states a name can be in and how state transitions work. Names are organized into namespaces. A namespace is the functional equivalent of a top-level domain in DNS—it defines the cost of names, and the name’s renewal rate. Like names, a namespace must be preordered and then registered.

## V. CONCLUSION

Namecoin, which is one of the oldest and most reputed cryptocurrency disappointed us under production environment. We got to know that how a single miner could have more than 51% hashing power which is a serious security issue and the network reliability was also not as good as Bitcoin. Moreover, even the best and reputed network like Namecoin are not prepared for the production environment. It is clear from the findings that Bitcoin is best when it comes to security.

We introduced Blockstack which is an open source blockchain based naming system which enables us to add new functionality without doing any modification in the underlying blockchain. With the new functionality of the Blockstack and several improvements it is easier to build decentralized services using PKI.

## ACKNOWLEDGMENT

Firstly we would like to thank the whole Blockstack community <http://blockstack.org> for helping us with the various discussions and allowing us to experiment their platform. Several Bitcoin developers like Peter Corallo and many more. Matt Todd helped us with the idea part of building Namecoin based system upon the Bitcoin network. Finally, we would like to thank several freelance developers for helping us with various issues.

## REFERENCES

- [1] Bitcoin hashrate. <https://blockchain.info/charts/hash-rate>.
- [2] Blockchain id format, ver 2. <https://github.com/blockstack/blockchain-id/wiki/Profile-Schema-v2>.
- [3] Crypto-currencies statistics – active nodes. <https://bitinfocharts.com>.
- [4] Cryptocurrency market cap, Jan. 2016. <http://www.coincap.io>.
- [5] Adam Back. Hashcash - A Denial of Service Counter-Measure. Tech report, 2002. <http://www.hashcash.org/papers/hashcash.pdf>.
- [6] Blockstack source code release v0.10, 2016. <http://github.com/blockstack/blockstack-server>.
- [7] J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten. Sok: Research perspectives and challenges for bitcoin and cryptocurrencies. In 2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015, pages 104–121, 2015.
- [8] Coindesk. State of bitcoin 2015: Ecosystem grows despite price decline, 2015. <http://coindesk.com/1tJDDvv>.
- [9] Coindesk. Bitcoin venture capital, Jan. 2016. <http://www.coindesk.com/bitcoin-venture-capital/>.
- [10] Counterparty protocol specifications. [http://counterparty.io/docs/protocol\\_specification/](http://counterparty.io/docs/protocol_specification/)
- [11] Filecoin: A Cryptocurrency Operated File Network. Tech report, 2014. <http://filecoin.io/filecoin.pdf>.
- [12] M. Gronager. Namecoin was stillborn, i had to switch off life-support, Oct 2013. <https://bitcointalk.org/index.php?topic=310954>.
- [13] D. Johnson, A. Menezes, and S. A. Vanstone. The elliptic curve digital signature algorithm (ecdsa). *Int. J. Inf. Sec.*, 1(1):36–63, 2001.
- [14] Juan Benet. IPFS - Content Addressed, Versioned, P2P File System. Draft, ipfs.io, 2015. <https://github.com/ipfs/papers>.
- [15] Jude Nelson and Larry Peterson. Syndicate: Virtual cloud storage through provider composition. In ACM BigSystem, 2014.
- [16] H. Kalodner, M. Carlsten, P. Ellenbogen, J. Bonneau, and A. Narayanan. An empirical study of Namecoin and lessons for decentralized namespace design. WEIS '15: Proceedings of the 14th Workshop on the Economics of Information Security, June 2015.
- [17] D. Kaminsky. Spelunking the triangle: Exploring aaron swartz's take on zookos triangle, Jan 2011. <http://dankaminsky.com/2011/01/13/spelunk-tri/>
- [18] J. A. Kroll, I. C. Davey, and E. W. Felten. The economics of bitcoin mining, or bitcoin in the presence of adversaries. In WEIS 2013.