Why Preventing a Cryptocurrency Exchange Heist Isn't Good Enough

Patrick McCorry[1], Malte Möser[2] and Syed Taha Ali[3]

[1] University College London, UK
p.mccorry@ucl.ac.uk,
[2] Princeton University, US
mmoeser@cs.princeton.edu,
[3] National University of Sciences and Technology, Pakistan
taha.ali@seecs.nust.edu.pk

Abstract. Cryptocurrency exchanges have a history of deploying poor security policies and it is claimed that over a third of exchanges were compromised by 2015. Once compromised, the attacker can copy the exchange's wallet (i.e. a set of cryptographic private keys) and appropriate all its coins. The largest heist so far occurred in February 2014 when Mt. Gox lost 850k bitcoins and unlike the conventional banking system, all theft transactions were irreversibly confirmed by the Bitcoin network. We observe that exchanges have adopted an overwhelmingly preventive approach to security which by itself has not yet proven to be sufficient. For example, two exchanges called NiceHash and YouBit collectively lost around 8.7k bitcoins in December 2017. Instead of preventing theft, we propose a reactive measure (inspired by Bitcoin vaults) which provides a fail-safe mechanism to detect the heist, freeze all withdrawals and allow an exchange to bring a trusted vault key online to recover from the compromise. In the event this trusted recovery key is also compromised, the exchange can deploy a nuclear option of destroying all coins.

1 Introduction

Cryptocurrencies have exploded into the mainstream over the last two years and now represent a thriving $700 billion ecosystem. Bitcoin, which was at parity with the US dollar in 2011, briefly traded for $19k per bitcoin in December, 2017, and is now recognised as legal tender in Japan and Germany. A Bitcoin futures contract has been formally launched at CME, the world's largest futures exchange. Amid all this positive press and mainstream recognition, however, cryptocurrencies continue to be dogged by the recurring scandal of hacked exchanges, resulting in billions of dollars in customer losses, and undermining user confidence in cryptocurrencies.

Cryptocurrency exchanges provide several valuable services for users. They serve as a convenient entry point for those wishing to purchase coins using conventional payment mechanisms. Customers can easily use the exchange platform to trade coins from one cryptocurrency to another. Most exchanges also provide

user wallets, thus enabling users to participate in transactions while sparing them the hassle of managing cryptographic keys. However, for all this convenience, this arrangment carries inherent counterparty risk: exchanges have full custody over customer coins and, due to the decentralized nature of cryptocurrencies, in case of theft, customer coins may be irretrievably lost.

As Moore et al. have documented, this risk is real: between 2009–2015 over a third of cryptocurrency exchanges were compromised and nearly half of all exchanges have simply disappeared [8]. The most well-known case is that of Mt. Gox, a Tokyo-based exchange, which at its peak was the world's largest Bitcoin exchange, handling 70% of all Bitcoin transactions, and shut down when $450 million worth of users' coins were stolen. This trend continues to this day. In December 2017, two exchanges were hacked: NiceHash lost 4.7k BTC (i.e. $45m US) [7] but remains operational, whereas YouBit lost 4k BTC (i.e. $39m US) [3], which represents 17% of its assets, and eventually declared bankruptcy.

A number of solutions have been proposed to reduce the risk for exchanges. Prominent proposals include the use of 'cold' wallets to store private keys where attackers cannot access them [31, 27, 26], hardware security modules to safeguard the hot wallet (i.e. private keys that are always on-line) [2] and the introduction of threshold signatures to distribute transaction authorization across multiple parties [1, 11]. However, in spite of these measures exchanges continue to fail. To further evaluate the impact of these failures, we provide a survey on large-scale heists and their cases in Section 2.

We argue that one major reason for this continued failure is that most security solutions employ a *proactive*, or preventive, approach. While this approach may be sufficient for conventional financial systems (where a centralized authority can – within certain limits – reverse transactions), it is clearly inadequate for the cryptocurrencies paradigm where damage can be done immediately and irreversibly. Instead, exchanges need to incorporate *reactive* mechanisms into their defense strategies that anticipate failure and allow them to recover *after* a successful attack.

To this end, we propose two reactive mechanisms to complement existing measures: our first solution allows the hot wallet to authorise time-delayed transfers and move coins between wallets within an exchange. This reduces the risk of theft by copying the cold storage's private keys as they may no longer be brought on-line periodically and the associated time-delay provides time for an exchange to detect abnormal transfers between their wallets. Our second proposal is a simple and intuitive failsafe mechanism (inspired by Möser et al.'s proposal of 'vault' transactions [23]) whereby exchanges may effectively detect a heist, freeze all time-delayed transfers, and recover by employing a trusted vault key. If the trusted vault is also compromised, the fail-deadly option can be used to destroy the stolen coins. Both solutions rely on the flexibility and expressiveness afforded by smart contracts and may be deployed on Ethereum without requiring any modification of the underlying platform.

2 Brief Survey of Exchanges Heists

We provide a brief survey on the causes behind large-scale exchange heists which includes the theft of wallets, insider threats and the rise/fall of distributed signing. As we will see, the focus of all security measures deployed by exchanges in response to heists are proactive (and preventive) in nature.

Theft of wallet (and address re-use) Mt. Gox remains the largest heist in the history of cryptocurrencies, with over 850k BTC stolen between 2011–2014. It was claimed by Mt. Gox that the loss of these coins were due to an underlying bug in Bitcoin called transaction malleability, but this was quickly debunked by Decker and Wattenhofer [10]. It was later claimed by WizSec [24] that MtGox's private keys were compromised in September 2011 and the company did not deploy any auditing mechanisms to detect the hack. The stolen set of keys were used to continuously steal new deposits as MtGox re-used Bitcoin addresses regularly and by mid 2013 over 630k BTC were stolen from the exchange. Remarkably, to support this claim, WizSec argues that evidence of the continuous theft can be extracted from transactions on the blockchain.

Reduce theft impact with hot and cold wallets To avoid significant losses as seen with Mt. Gox, many companies incorporate cold and hot wallets. All coins are sent to the exchange's cold wallet and when necessary coins are manually transferred from the cold to hot wallet. If an exchange's server is compromised, then only coins in the hot wallet can be stolen by the thief and thus an exchange can determine the number of coins it is willing to risk. While these hacks are less severe, high-impact heists continued, including 24k BTC stolen from BitFloor in May 2012 [17], 19k BTC stolen from Bitstamp [18] in June 2015 and more recently in December 2017 where 4.7k BTC were stolen from NiceHash and 4k from YouBit. Both BitFloor and YouBit declared bankruptcy as a result of these hacks, whereas BitFloor and Bitstamp remain operational.

Insider threat While keeping keys offline makes stealing coins harder, the compromise of cold wallets is not infeasible as witnessed in Feburary 2015, when an exchange called BTER claims to have lost 7,170 BTC from their cold wallet and it is rumoured this heist was due to an insider [14]. Another exchange, Shapeshift has provided a post-mortem [28] on three seperate compromises that led to a heist of 315 BTC. The report mentions the installation of a rootkit and that the thief purchased an SSH Key and the username/password of ShapeShift HQ's router from a former employee. Note due to the nature of ShapeShift, most customer funds were not at risk and it was in fact ShapeShift's own coins that were stolen. We highlight this is because ShapeShift only has temporary custody over coins in order to facilitate an exchange.[4]

[4] ShapeShift is a match-making exchange and sends all customers coins in the corresponding cryptocurrency once the exchange is complete. For example, the customer may send ShapeShift bitcoins and shortly afterwards ShapeShift will send the customer ether.

Compromised shared web hosting Linode is a web hosting company offering virtual server rentals and several promiment Bitcoin exchanges/wealthy community members used Linode to store their hot wallets. In June 2011, an attacker compromised Linode and targeted the virtual services that stored the hot wallets. Unfortunately this led to the theft of at least 46k bitcoins and the exact amount remains unknown. [12] The victims included Bitcoinia at approximately 43k BTC and Bitcoin.cx at 3k BTC, as well as bitcoin developer Gavin Andresen who also lost 5k BTC. Note after the Meltdown/Spectre disclosure [19, 16] in January 2018, companies such as Coinbase highlighted that they still rely on shared web hosting for non-sensitive workloads:

Coinbase runs in Amazon Web Services (AWS) and our general security posture is one of extreme caution. Sensitive workloads, especially where key handling is involved, run on Dedicated Instances (instead of shared hardware). Where we do run on shared hardware, we make it more difficult to accurately target one of our systems by rapidly cycling through instances in AWS.
Philip Martin, Director of Security at Coinbase [21]

Rise and Fall of Multisig Over the course of 2011 and 2012, Andresen incorporated threshold signatures (i.e. multisig) into Bitcoin [1] which was slowly adopted by the community and we highlight that in Feburary 2018 there is approximately 3.6m bitcoins (and 11.8m outputs) that rely on multi-sig.[5] In the aftermath of both hacks of Linode and Bitfloor, Andresen argued that the new multisig feature could reduce the likelihood of heists in exchanges as it requires the attacker to compromise a threshold of machines instead of a single machine. As well, in early 2014, the security company BitGo further declared that the rise of multisig would remove the need for cold wallets (i.e. offline keys) altogether:

"We (BitGo) believe it's time we come together as an industry to end the cold storage ice age and adopt multi-sig."
Will O'Brien, co-founder at BitGo [25]

However, the use of multisig is no silver bullet in itself, and this was proven by another large heist at Bifinex with 119,756 BTC stolen. Bitcoin exchange Bitfinex partnered with BitGo and used them as a third party escrow to audit/approve customer withdrawals. Furthermore it appears that Bitfinex decided not to use cold wallets (and opted for a third-party auditor instead) [6] to avil a statutory exemption to the Commodities and Exchange Act. It was later revealed that BitGo had a special configuration for BitFinex [4] and while a final report on the heist remains unpublished, it appears BitGo simply authorised all transaction requests from BitFinex. This highlights that while the concept of using threshold signatures is attractive, it does not guarantee that the authority to authorise transactions is in fact distributed.

[5] p2sh.info tracks the number of *pay-to-script-hash* outputs which are mostly multi-sig scripts.

Impact of proactive security All security measures deployed so far by cryptocurrency exchanges are proactive in nature with the goal to prevent a heist. The above survey highlights that proactive security measures have evidently reduced the impact of heists, but unfortunately they cannot prevent a heist. Fundamentally, the issue is that once the corresponding private keys are stolen, there is little an exchange can do in order to stop the heist due to the irreversible nature of the blockchain. In the next section, we highlight a potential reactive solution that requires a time-delay on all withdrawals which provides a grace period for an exchange to react and cancel the heist.

3 Reactive Security Measures for Cryptocurrencies

We propose that cryptocurrency exchanges should deploy reactive security measures that allow them to respond in the event of a heist. Our solution has two components which includes time-delayed (and revocable) payments and time-delayed access control of hot/cold wallets.

A revocable and time-delayed payment was first proposed for Bitcoin as a *vault transaction* [23], where a payment is initiated by a key that is held in a *hot wallet*, but can later be revoked (i.e. recovered) within a certain time frame by a key held securely in cold storage [23]. We note that such a mechanism is currently not possible in all cryptocurrencies: while it is trivial to deploy using Ethereum's smart contract language (which we show later in this section), Bitcoin's scripting language requires support for covenants (cf. [23]).

We argue that time-delayed and revocable payments cannot only be used to secure funds, but also to simplify access control mechanisms. For example, an exchange contract may forsee the funds in its hot wallet running low and authorise a time-delayed transfer of funds from cold to hot storage using its hot keys. The cold storage signing keys can stay offline, they are only brought online if the coins need to be transferred immediately or the payment needs to be revoked. In the following, we describe an advanced vault design that combines multiple such mechanisms.

3.1 Time-Delayed Exchange Vaults

We propose Time-Delayed Exchange Vaults, inspired by Möser et al's Bitcoin vaults [23], which combine multiple proactive and reactive mechanisms: they allow customers to deposit coins into cold storage, and exchanges to perform time-delayed transfers of coins from cold to hot storage using their online keys and authorise time-delayed customer withdrawals. If an exchange is compromised, the exchange can lock down the vault using an online key (including the compromised key) and effectively freeze all withdrawals. This provides time for the exchange to bring a trusted vault key online and recover from the compromise. We outline the proposed protocol before presenting a proof of concept implementation.

Proposed Protocol We present how to establish the contract, the three types of keys involved in access control of the coins, the time-delay involved in transferring coins from cold to hot storage and withdrawing coins, and finally the lock down and recovery process.

Contract set up and access control There are three sets of keys that must be set when the contract is established in order to facilitate access control:

- *Hot keys:* Always online to authorise customer withdrawals,
- *Cold keys:* Periodically online to transfer coins from cold to hot wallet,
- *Vault key:* Only used to re-issue hot/cold keys and restart withdrawals, or if compromised can destroy all coins in the wallet.

The exchange can store a list of each key type and require a threshold of k out of n keys to sign and authorise moving coins. Two timers $t_{coldtransfer}$ and $t_{withdrawal}$ are required to self-enforce the time-delay, with the former controlling the time required to transfer coins from cold to hot storage and the latter delaying customer withdrawals.

Transferring coins There are two types of messages that can be signed by an exchange to authorise moving coins: *cold transfer request* and *withdrawal request*. The former moves coins from cold to hot storage which is instant if signed by a threshold of cold keys or incorporates a time-delay of $t_{coldtransfer}$ if signed by a threshold of hot keys. The latter associates a customer with coins from the hot wallet which can be withdrawn after $t_{withdrawal}$ and must be signed by a threshold of hot keys.

Vault lock down If an exchange has detected a heist, it is responsible for signing a *lock down* message from any key registered in the contract. This message disables all functionality within the contract except for the recovery procedures which can only be accessed using the trusted recovery key. The exchange can set a new list of hot/cold keys and a new recovery key before canceling all withdrawal requests, transferring all coins to cold storage and re-enabling the functionality within the contract. If the trusted recovery key is also compromised, then the exchange can deploy a nuclear option of signing a *destroy* message using the recovery key. This disables all functionality within the contract and removes access to the coins forever. As a result neither the thief or exchange will control the coins.

Future work We plan to investigate whether the contract can self-enforce rate limits and detect abnormal spending behaviour in a cost-efficient manner. If so the contract can automatically increase the time-delay. We also need to identify whether an on-chain rate limit is desirable or not. An in-progress implementation is available.[6]

[6] Link to in progress implementation: https://pastebin.com/KiPGB23k

4 Discussion

Perception of time delayed withdrawals. It is common for users to publicly complain about slow withdrawals of coins/fiat currency. Recent complaints in 2017 include users waiting up to 30 hours for withdrawals from Coinbase [9] and the need for Bitfinex to manually verify withdrawals after denial of service attacks between the 4-5th December [5]. While the time delay for withdrawals in our proposed solution is publicly verifiable by the customer, it may be undesirable to further increase the delay for customers receiving their deposits.

Financial privacy for the exchange The benefit of the Bitcoin vault approach is that the underlying nature of Bitcoin's transaction design makes it difficult (but not impossible [22]) to identify the coins held by an exchange as all deposits are distinct sets of coins. In contrast, our smart contract solution provides little privacy as all coins are stored in a single location. This not only leaks information about an exchange's assets and earnings, but also about internal structures and processes as all keys and access rights have to be specified in the contract.

Risks of smart contract wallets. There is a growing list of smart contracts that have resulted in the theft or loss of coins due to subtle bugs in their code. Two prominent examples include TheDAO [13] where a thief exploited a bug to steal over 3.5m ether and the community was required to co-ordinate a hard-fork to stop the theft, and the Parity wallet [30] bug that allowed a novice user to kill a central smart contract and effectively freeze around 519k ether. We highlight that while there is an accumulating list of methods to formally verify the correctness of a smart contract [20, 32, 15], it remains to be seen whether an exchange is willing to risk all their assets to a single contract.

Bitcoin vaults vs contract vaults. In vault transactions [23], due to the complexity of Bitcoin's transaction design, each set of stolen coins requires a new signature from the vault key for recovery. Signing multiple (and potentially large) transactions is problematic due to the rise of congested blocks in Bitcoin and significant fees (i.e. $20 or more for a small transaction). Exchanges often have coins spread across a large set of outputs (e.g., Coinbase has more than 1.5 million sets of spendable coins for just 265 BTC [29]), requiring a significant amount of signatures from the vault key. While batching funds is conceivable, the fact that each set of coins needs to store the entire list of access options is a significant drawback of Bitcoin.

Our proposed solution demonstrates that when designing vault transactions for smart contract-enabled cryptocurrencies like Ethereum, the account-based ledger (unlike Bitcoin's complex unspent transaction output ledger) allows a single transaction to lock down and cancel all withdrawals. The expressiveness of smart contracts allows any registered key (and not just the trusted vault key) to perform the lock down which allows an exchange to quickly react in the event of a heist. Finally, more granular access control can be incorporated in the contract to allow the hot keys to transfer coins from cold-storage with

a time-delay such that the cold keys no longer need to periodically be brought online and thus mitigate the risk of their theft.

Fail-deadly mechanism. A nuclear option to destroy all coins is a strong deterrent against attacks, as the attacker knows beforehand that even if she can compromise keys, all funds will be destroyed before they can be moved outside of the contract. However, it is an open question whether companies would actually make use of such a feature. While it reduce the set of possible attackers (to those that gain indirectly by the exchange having to close), it would ruin the company, could confuse or deter users and might raise regulatory concerns.

References

1. G. Andresen. Pay to Script Hash. *Bitcoin Github*, Jan. 2012.
2. N. Bacca. How to properly secure cryptocurrencies exchanges. *Ledger Company*, Aug. 2016.
3. BBC. Bitcoin exchange Youbit shuts after second hack attack. *BBC*, Dec. 2017.
4. M. Belshe. Bitfinex Breach Update. *BitGo Blog*, Aug. 2016.
5. Bitfinex. Explanation . *Reddit Post*, Dec. 2017.
6. J. Brito. What does the CFTC have to do with the Bitfinex hack? *Coin Center*, Aug. 2016.
7. R. Browne. More than $60 million worth of bitcoin potentially stolen after hack on cryptocurrency site. *CNBC*, Dec. 2017.
8. G. Chavez-Dreyfuss. Cyber threat grows for bitcoin exchanges. *Reuters*, Aug. 2016.
9. Dariusz. Slow Withdrawals Leave Coinbase Users Annoyed. *TheMerkle*, 2017.
10. C. Decker and R. Wattenhofer. Bitcoin transaction malleability and mtgox. In *European Symposium on Research in Computer Security*, pages 313–326. Springer, 2014.
11. R. Gennaro, S. Goldfeder, and A. Narayanan. Threshold-optimal dsa/ecdsa signatures and an application to bitcoin wallet security. In *International Conference on Applied Cryptography and Network Security*, pages 156–174. Springer, 2016.
12. D. Goodin. Bitcoins worth $228,000 stolen from customers of hacked Webhost. *Arstechnica*, Feb. 2012.
13. R. Hanson. A $50 Million Hack Just Showed That the DAO Was All Too Human. *Wired*, June 2016.
14. S. Higgins. BTER Claims $1.75 Million in Bitcoin Stolen in Cold Wallet Hack. *Coindesk*, Feb. 2015.
15. E. Hildenbrandt, M. Saxena, X. Zhu, N. Rodrigues, P. Daian, D. Guth, and G. Rosu. Kevm: A complete semantics of the ethereum virtual machine. Technical report, 2017.
16. P. Kocher, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom. Spectre attacks: Exploiting speculative execution. *arXiv preprint arXiv:1801.01203*, 2018.
17. T. B. Lee. Hacker steals $250k in Bitcoins from online exchange Bitfloor. *Arstechnica*, May 2012.
18. R. Lemos. Bitcoin exchange Bitstamp claims hack siphoned up to $5.2 million. *Arstechnica*, 2015.
19. M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg. Meltdown. *arXiv preprint arXiv:1801.01207*, 2018.
20. L. Luu, D.-H. Chu, H. Olickel, P. Saxena, and A. Hobor. Making smart contracts smarter. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 254–269. ACM, 2016.
21. P. Martin. Coinbase Accused of Technical Incompetence After Hoarding Millions of UTXOs. *Coinbase*, Jan. 2018.
22. S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage. A fistful of bitcoins: characterizing payments among men with no names. In *Proceedings of the 2013 conference on Internet measurement conference*, pages 127–140. ACM, 2013.

23. M. Möser, I. Eyal, and E. G. Sirer. Bitcoin covenants. In *International Conference on Financial Cryptography and Data Security*, pages 126–141. Springer, 2016.
24. K. Nilsson. Breaking open the MtGox case, part 1. *Wizsec*, July 2017.
25. W. O'Brien. It's Time to End the Cold Storage Ice Age and Adopt Multi-Sig. *BitGo*, Sept. 2014.
26. M. Palatinus and P. Rusnak. Multi-Account Hierarchy for Deterministic Wallets. *Bitcoin Github*, Apr. 2014.
27. M. Palatinus, P. Rusnak, A. Voisine, and S. Bowe. Mnemonic code for generating deterministic keys. *Bitcoin Github*, Sept. 2013.
28. M. Perkin. Bitfinex Breach Update. *LedgerLabs*, 2016.
29. K. Sedgwick. Coinbase Accused of Technical Incompetence After Hoarding Millions of UTXOs. *bitcoin.com*, Dec. 2017.
30. L. Shen. Millions of Dollars Worth of Ethereum Got Locked Up. Here's Why. *Fortune*, Nov. 2017.
31. P. Wuille. Hierarchical Deterministic Wallets. *Bitcoin Github*, Feb. 2012.
32. E. Zurich. Formal Verification of Ethereum Smart Contracts. *Securify*, Jan. 2017.