Cyber Security Threats to Bitcoin Exchanges
Adversary Exploitation and Laundering Techniques

Oosthoek, Kris; Doerr, Christian

DOI
10.1109/TNSM.2020.3046145

Publication date
2021

Document Version
Final published version

Published in
IEEE Transactions on Network and Service Management

Citation (APA)
Oosthoek, K., & Doerr, C. (2021). Cyber Security Threats to Bitcoin Exchanges: Adversary Exploitation and Laundering Techniques. *IEEE Transactions on Network and Service Management*, *18*(2), 1616-1628. Article 9300238. https://doi.org/10.1109/TNSM.2020.3046145

Cyber Security Threats to Bitcoin Exchanges: Adversary Exploitation and Laundering Techniques

Kris Oosthoek, *Member, IEEE*, and Christian Doerr, *Member, IEEE*

***Abstract*—Bitcoin is gaining traction as an alternative store of value. Its market capitalization transcends all other cryptocurrencies in the market. But its high monetary value also makes it an attractive target to cyber criminal actors. Hacking campaigns usually target an ecosystem's weakest points. In Bitcoin, the exchange platforms are one of them. Each exchange breach is a threat not only to direct victims, but to the credibility of Bitcoin's entire ecosystem. Based on an extensive analysis of 36 breaches of Bitcoin exchanges, we show the attack patterns used to exploit Bitcoin exchange platforms using an industry standard for reporting intelligence on cyber security breaches. Based on this we are able to provide an overview of the most common attack vectors, showing that all except three hacks were possible due to relatively lax security. We show that while the security regimen of Bitcoin exchanges is subpar compared to other financial service providers, the use of stolen credentials, which does not require any hacking, is decreasing. We also show that the amount of BTC taken during a breach is decreasing, as well as the exchanges that terminate after being breached. Furthermore we show that overall security posture has improved, but still has major flaws. To discover adversarial methods post-breach, we have analyzed two cases of BTC laundering. Through this analysis we provide insight into how exchange platforms with lax cyber security even further increase the intermediary risk introduced by them into the Bitcoin ecosystem.**

***Index Terms*—Bitcoin, cryptocurrency exchanges, cyber security, cyber threat intelligence, attacks, vulnerabilities, forensics.**

I. Introduction

With an average market capitalization of 136 billion USD over the last two years [1], Bitcoin transcends all other currencies in the cryptocurrency market space. Similar to other currencies, security is a critical property in securing its role as a store of value, unit of account and means of exchange. Owners have to be confident that they won't lose their funds or they will withdraw them. While the developers of Bitcoin's reference implementation, Bitcoin Core, acknowledge that certain attack vectors exist [2], their probability is low as long as honest Bitcoin nodes together control more processing power than any group of attacker nodes [3]. Due to its implementation of a stack of cryptographic technologies, Bitcoin is a safe and reliable digital currency in its core. This article will not review fundamental attacks on Bitcoin's distributed ledger technology, but considers another type of attack that has been proven most lucrative and continues to be.

A high value asset makes a high value attack target. The security of Bitcoin is also dependent on the ecosystem that has emerged around it. This consists of exchange platforms, payment service providers, wallet providers, mining pools and other intermediaries. Each of these is part of a fabric spun around Bitcoin which unlocks its potential to a broader user base, but consequently introduces additional threat vectors.

Cyber criminal actors generally target the weakest points in the ecosystem. In the Bitcoin ecosystem, centralized exchanges make up a large part of these. These act as a broker, allowing users to sell cryptocurrencies for fiat currency (legal tender) or to exchange the latter for cryptocurrency against a commission. Attacks on their platforms are feasible because in contrast to conventional stock exchanges, they do store currencies traded or exchanged by their clients.

This contradicts the original Bitcoin proposition as a decentralized currency, in which ownership depends on knowledge of the public-private key pair. The keys are the money: "not your keys, not your Bitcoin" [4]. However many owners deposit their Bitcoin with the exchange, which acts as a custodian. Although storing Bitcoin with an intermediary is a compelling offer for users such as active traders requiring quick and easy access to their funds, it creates a false sense of security to users less informed about the security aspects of Bitcoin ownership. Control of funds and thus the exercise of ownership is outsourced to or centralized at the exchange. While legal ownership is non-transferable, the public-private key pair that implies ownership of BTC remains with the exchange. According to recent reporting, the biggest exchange holds 966,230 BTC in custody, worth 7.19 billion USD at the moment of writing [5]. Exchanges must implement bank-level security to avoid successful cyber attacks and to safeguard funds, but have failed to do so as proven by the breaches in our analysis.

Bitcoin that are not being actively traded should be stored in cold storage. The hardware wallet is the best-known and most end-user friendly solution for cold storage. While the transfer of user funds to cold storage is a security best practice for exchanges, they regularly have funds available in hot storage in order to provide for quick exchange or withdrawal by legitimate users. With hot wallets being directly connected to the

Manuscript received June 28, 2020; revised October 9, 2020; accepted November 11, 2020. Date of publication December 21, 2020; date of current version June 10, 2021. The associate editor coordinating the review of this article and approving it for publication was A. Veneris. *(Corresponding author: Kris Oosthoek.)*

Kris Oosthoek is with the Cyber Threat Intelligence Lab, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: k.oosthoek@tudelft.nl).

Christian Doerr is with the Cyber Threat Intelligence Lab, Hasso Plattner Institute for Digital Engineering, University of Potsdam, 14469 Potsdam, Germany (e-mail: christian.doerr@hpi.de).

Digital Object Identifier 10.1109/TNSM.2020.3046145

Internet and running as an ongoing process to rapidly meet liquidity requirements, they introduce the risk of exchange platforms losing BTC through exploitation of unknown vulnerabilities in their infrastructure. Cyber security is not top of mind in the development process of many start-up technology companies, which most centralized exchange platforms are. This has resulted in frequent reports of client funds getting lost due to breaches. According to a March 2019 report from the United Nations Security Council, cryptocurrency exchanges are even targeted by sophisticated nation-state hacking groups in order to fund nuclear weapons programs [6].

With BTC market capitalization growing over time, attention to exchange platform security must grow in importance. Each incident potentially not only has a monetary impact, but potentially affects Bitcoin's credibility as a monetary asset.

This article is an invited extended version of a paper presented at the 2020 IEEE International Conference on Blockchain and Cryptocurrency [7]. For this article we have extended our work with an analysis of the security posture of the exchange platforms in our dataset that are still active. In addition to that we also analyze how the Lazarus group and the actor group behind the Bitfinex breach are laundering stolen BTC. Even during our analysis, more than 4 years after the Bitfinex breach, transactions with wallets linked to the hack were still observed.

Our systematic study of Bitcoin exchange breaches provides the following take-aways and contributions.

- We show that most Bitcoin exchanges were breached through relatively straightforward attack vectors.
- We found that while attack vectors overlap with breaches of other financial service providers, the actual exfiltration of funds is unique to Bitcoin exchanges.
- We demonstrate that over recent years the sophistication of the vectors used to breach exchanges has increased.
- We found that while the amount of BTC stolen per breach tends to decrease, the USD yield is higher due to an increased BTC-USD exchange rate.
- We demonstrate that the age of breached exchanges has increased in recent years.
- We show that over recent years more exchanges tend to survive after a breach, but details on the attack vector used are shared decreasingly.
- We demonstrate that while security has improved over time, some platforms still have relative lax Web security when held against standards such as OWASP.
- We provide insight into adversary methods to launder stolen BTC through the blockchain and that the conversion of BTC to fiat money has become more complex.

The remainder of this article is structured as follows: Section II provides an overview of related work on Bitcoin exchange security. Section III provides an overview of the Cyber Threat Intelligence field and the Vocabulary for Event Recording and Incident Sharing. Section IV describes the methodology of our analysis. Section V presents the results from our analysis. Section VI provides an overview of the Web security posture of exchange platforms. Section VII describes adversary laundering techniques post-breach. Section VIII outlines the limitations of our research. Section IX summarizes our findings.

II. Related Work

Several authors have focused on theoretical attacks on the Bitcoin network. The extensive research of Conti *et al.* has delivered a reference article on security and privacy concerns regarding Bitcoin. Their article focuses on various attack types such as double spending, Finney, brute force, Vector 76 and Goldfinger attacks [8]. They also cover the various countermeasures for these attacks. Lim *et al.* have also focused on security threats to Bitcoin such as DDoS attacks against exchanges, Bitcoin mining malware and extortion [9]. Feder *et al.* have looked at the impact of DDoS attacks on the now defunct Mt. Gox exchange. They found that on days following DDoS attacks, trading volume significantly decreased, specifically caused by a drop in large volume trades [10].

With regards to the risks introduced by Bitcoin exchanges in particular, Moore *et al.* have looked at various risk factors that have influenced the closure of Bitcoin exchanges between 2010 and 2015 [11]. They found that nearly half of the exchanges in their dataset have closed due to fraud attempts and security breaches. While they mention Bitcoin exchanges as the scope of their dataset, their analysis does also include services that did not support Bitcoin, e.g., Ripple. Their analysis is particularly useful, as they have analyzed the relationship between the presence of security-related features such as multiple factor authentication, bug bounties and exchange closure. Their dataset is however a bit outdated. They also have more of an economic focus on the exchange ecosystem and focus less on the actual security problems through which breaches have occurred. With regards to the cyber threats to Bitcoin exchanges specifically, several online resources that provide unstructured overviews of breaches exist [12], [13].

Various authors have focused on the topics Bitcoin exchanges and Bitcoin security independently of each other. However we did not find any peer-reviewed contributions on the cyber security of exchange platforms. As far as we are aware, any significant academic analysis of a corpus of Bitcoin exchange breaches has not been performed, which we deem the main contribution of our research.

III. Cyber Threat Intelligence

Cyber Threat Intelligence (CTI) is an umbrella term for the analysis of cyber security breaches and their tools, tactics and procedures (TTPs). It aims to provide actionable information to drive cyber security decision-making in order to avoid getting attacked with TTPs that were already disclosed. As there are many cyber threats around which are not all relevant to each organization, CTI aims to provide an understanding of the threats relevant to an organization and its assets. In this article we focus on cyber threats to Bitcoin exchanges.

The CTI process strives to gain an information advantage on adversarial events to an organization's information systems. Threats are real if they are able to successfully exploit a vulnerability, leading to a normally negative real-world impact. The

malicious actor needs to have the capability and opportunity to exploit that vulnerability and the intent to do bad things. Commonly heard attack vectors like ransomware, Denial of Service attacks, SQL injection and phishing can have a different impact for each individual organization as these depend on particularities specific to their technical environment.

Several frameworks to understand cyber threats in context exist, such as STRIDE, CAPEC, ATT&CK and VERIS. They each have their own distinct use case. STRIDE, a mnemonic for Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege, is useful to understand how threats can impact an information system in several ways [14]. CAPEC is useful for analysis of software exploit methods [15], whereas ATT&CK is proven to be useful to inform security analytics and understand malware trends [16]. For our analysis we have used VERIS, which is primarily useful for post-breach assessments.

A. Vocabulary for Event Recording and Incident Sharing (VERIS)

The Vocabulary for Event Recording and Incident Sharing (VERIS) is a CTI standard open-sourced by Verizon [17]. Of all efforts that exist to systematize the conversation on cyber security and exchange information, VERIS has become the industry standard for strategic CTI. It is targeted towards strategic CTI as it is meant for reporting that informs strategic, longer-term decision making to prioritize security investments based on risk appetite. Four indicators present in every cyber security incident form the basis of how VERIS is structured: the *Action* used to breach the asset, the *Actor* who breached the asset, the compromised *Asset*, the security *Attribute* (confidentiality, integrity or availability) that was affected.

Analysis of TTPs in a set of cyber security breaches can provide an understanding of how attackers target an industry such as Bitcoin exchanges. VERIS provides for structured analysis as it translates the narrative of individual incidents into a structured form. The most well-known example of VERIS in use on a multiple-industry level is the annual Verizon Data Breach Report (DBIR), which has become the industry-standard reference for intelligence on developments in the cyber threat landscape [18]. Breach data from a large body of industry and public sector organizations is used as the source for the report. Publicly disclosed breaches are also recorded in the VERIS Community Database, available on Github [19].

IV. Methodology

In this section the methods applied in the collection of our dataset of Bitcoin exchange breaches and their successive classification using VERIS are discussed. In this article we use *Bitcoin* to refer to the Bitcoin distributed ledger and technology stack. *BTC* is used to designate units of account, e.g., when referring to the amount stolen in a particular breach.

A. Data Collection

1) Exchange Breaches: We have gathered our dataset in November 2019 using Google Custom Search JSON API queries for *bitcoin exchange breaches* and *bitcoin exchange thefts*. Based on word frequency analysis, we identified 36 incidents of breached exchange platforms, which were cross-checked against media reporting. Our analysis is concentrated on Bitcoin exchanges. Based on our criteria we only included technical *security* breaches of exchanges which focus either on *Bitcoin* trading exclusively or combined with other cryptocurrencies. Exchanges not supporting Bitcoin and exchanges without reported breaches are not included in our dataset. In some breaches of multi-currency platforms, other currencies than Bitcoin were stolen as well. In those cases we include the amount of BTC stolen according to official reports. Our dataset does not include any decentralized exchanges for peer-to-peer trading, hash-power marketplaces or online wallet services.

2) Financial Services Breaches: In order to compare Bitcoin exchange breaches with breaches of other financial service providers such as banks, we have used the VERIS Community Database (VCDB). This public dataset includes VERIS-formatted, annotated reports of publicly disclosed cyber security breaches in various industries. It is audited by the VERIS Risk team at Verizon and also used as input for their annual report. At the time of writing the database includes 8346 incidents, updated daily. Based on the VERIS taxonomy, it allows to filter for data breaches that occurred with organizations offering financial services. We have used the JSON objects of *validated* incidents, which are manually checked for validity by Verizon [20]. The VCDB captures incidents recorded from 2012 and is thus aligned with the time period covered by our dataset of exchange breaches, allowing for a uniform comparison. We have checked whether incidents from our exchange breach dataset are recorded in VCDB, however no overlap existed.

3) Trade Volumes: We have gathered data on daily exchange trade volumes from a public API offered by CoinGecko [21]. While CoinGecko is one of the few overview websites that normalizes data in order to account for exchanges reporting fake volume, currently no publicly available dataset exists that fully accounts for exchanges reporting fake volume data. This is a known problem of the exchange ecosystem [5]. For this reason, this data was only used to analyze post-breach impact as reported by exchanges in Section VI-B5.

B. Classification of Breaches

For our analysis we have focused on the Action category of VERIS. The VERIS taxonomy also has an Actor category, but rarely are breaches of Bitcoin exchange platforms attributed to designated actor groups. The Asset category is not used because for each exchange breach, the Server asset would qualify as the platforms in our dataset are online outlets exclusively. In case of another financial service provider like a bank, the breached asset can also be an ATM for example. The Attribute category affected would always be Confidentiality and Integrity, as we have not recorded Denial of Service attacks affecting platform availability.

The analysis of Bitcoin exchange breaches has proven to be onerous as the sharing of information tends to be quite scarce and is getting even more scarce over recent years. For

TABLE I
BITCOIN EXCHANGE BREACHES

Launch	Breach	Exchange	BTC	USD	Action	Variety	Attack method	Closed
2010-07	2011-06	Mt. Gox	2,500	40,250	Hacking	Use of stolen creds	admin account breach [26]	yes*
2010-04	2011-08	MyBitcoin	12,500	102,500	Hacking	Abuse of functionality	programming error [27]	yes
2011-09	2012-03	Bitcoinica	43,554	213,415	Hacking	Exploit vuln	Linode breach [28]	yes*
2011-09	2012-05	Bitcoinica	18,547	96,444	Hacking	Abuse of functionality	Rackspace PW recovery [29]	yes
2011-04	2012-05	BitMarket.eu	19,980	103,896	Hacking	Use of stolen creds	SSH account hacked [30]	yes
2012-02	2012-09	Bitfloor	24,086	298,666	Hacking	Use of stolen creds	unencrypted wallet backup [31]	yes
2011-08	2013-03	BitInstant	999	92,907	Social	Elicitation	DNS hijack, registrar social eng. [32]	yes
2011-01	2013-03	Mercado Bitcoin	4000	372,000	Hacking	Abuse of functionality	coupon functionality hijacked [33]	no
2010-12	2013-04	Bitcoin Central	'few 100'	-	Hacking	Abuse of functionality	account takeover, OVH breach [34]	no
2011-10	2013-05	Vircurex	1,454	187,275	Hacking	Use of stolen creds	PW reset, cloud host social eng. [35]	yes
2011-07	2013-11	Bitcash.cz	485	584,765	Hacking	Unknown	*web interface compromised* [36]	yes
2013-01	2013-11	Bidextreme.pl	-	-	Hacking	Unknown	[37]	yes
2010-07	2014-02	Mt. Gox	850,000	487,815,000	-	-	insider involvement [38]	yes
2014-01	2014-03	Poloniex	97	43,136	Hacking	Abuse of functionality	race condition [39]	no
2013-03	2014-07	Cryptsy	10,000	5,895,000	Hacking	Use of backdoor or C2	backdoor in dependency [40]	yes
2011-01	2015-01	Bitstamp	18,866	4,122,221	Malware	Downloader	sophisticated malware attack [41]	no
2013-06	2015-01	796 Exchange	1,000	218,500	Hacking	Unknown	compromised "certain weakness" [42]	no
2013-01	2015-02	BTER	7,170	1,821,897	Hacking	Unknown	breach of cold wallet [41]	yes
2011-06	2015-02	Cavirtex	-	-	Hacking	Use of stolen creds	PW hashes, 2FA secrets exposed [43]	no
2014-10	2015-02	Excoin	-	-	Hacking	Unknown	[44]	yes
2014-02	2015-03	AllCrypt	40	9,764	Hacking	Brute force	bruteforced tech staff email [45]	yes
2014-01	2015-03	Cryptoine	6	1,465	Malware	Exploit vuln	race condition in trading engine [46]	yes
2013-01	2016-03	Cointrader	-	-	-	-	[47]	yes
2013-07	2016-04	BitQuick	-	-	Hacking	SQLi	upload feature SQL injection [48]	no
2014-07	2016-04	ShapeShift.io	315	141,278	Hacking	Use of stolen creds	insider involvement [49]	no
2013-07	2016-05	Gatecoin	250	132,225	Hacking	Unknown	multisig cold wallet [50]	yes
2012-01	2016-08	Bitfinex	120,000	68,868,000	Hacking	Unknown	[51]	no
2012-01	2016-11	Bitcurex	2,300	1,707,750	Hacking	Unknown	*API signing key exploit* [51]	yes
2013-01	2017-04	Yapizon	3,816	5,158,850	Hacking	Unknown	[52]	no
2013-06	2018-04	Coinsecure	438	4,049,354	Misuse	Privilege Abuse	*insider*, cold storage exposed [53]	yes
2014-06	2018-09	Zaif	5,966	39,585,603	Hacking	Unknown	3 hot wallets hacked [54], [55]	no
2018-05	2018-10	Maplechange	8	50,927	Malware	Exploit vuln	race condition, *exit scam* [56]	yes
2013-07	2018-11	Gate.io	-	-	Malware	Exploit vuln	supply chain attack [57]	no
2017-11	2019-03	DragonEx	135	553,811	Hacking	Unknown	[58]	no
2017-07	2019-05	Binance	7,000	59,908,100	Hacking	Unknown	API keys and 2FA secrets [59]	no
2016-01	2019-07	BitPoint	1,225	12,350,450	Hacking	Unknown	[60]	no

* the asterisk in the Closed column indicates closure after a subsequent breach

our analysis, security breaches were included in our dataset according to the criteria by Verizon. The entry must be a confirmed security incident, with a loss of confidentiality, integrity, or availability [18]. In the case of our analysis, we also chose to only include officially disclosed breaches, meaning they were announced through official communication channels maintained by the particular exchange. Press releases, but also messages from the official Twitter channel or posts on *bitcointalk.org* from confirmed accounts of exchange staff. We provide references to each source.

Breaches were classified to a threat action category and variety according to the information we had on the *initial* point and means of entry, as this provides for an overview of the attack surfaces of Bitcoin exchanges. As in some cases official sources only reported a successful hacking attempt and lack further detail, we have not identified the sub-variety of Hacking to stay close to the official incident report. Also, in some cases the amount of BTC stolen is not reported. In other cases, a cyber security breach is the official account, but heavy rumors about a cover-up such as an exit scam exist. Because we want our analysis to be a valid but accurate reflection of the current state of the Bitcoin exchange ecosystem, we have included this in italic in the Attack Method column of Table I. We refer to the URLs used for the coding of each breach.

V. ANALYSIS OF BITCOIN EXCHANGE BREACHES

In this section we will discuss the observations from our dataset of Bitcoin exchange breaches. We have analyzed 36 incidents of breaches, through which at least 1,156,399 BTC were stolen from their legitimate owners. Table I and Figure 1 provide additional insight into the dataset and our analysis.

Figure 1 is a bubble chart representation of the dataset. Exchange breaches are plotted on the X axis by year of compromise. The Y axis indicates the age of the exchange at the time of the incident. Each bubble represents a breach, whereas the line color represents the attack vector used and the bubble diameter the amount of BTC stolen. Figure 1 shows two interesting patterns in particular. The amount of BTC stolen in breaches has decreased in recent years. It also shows that the age at which an exchange gets breached has increased. Furthermore the figure shows that the TTPs deployed in breaches of exchange platforms have developed

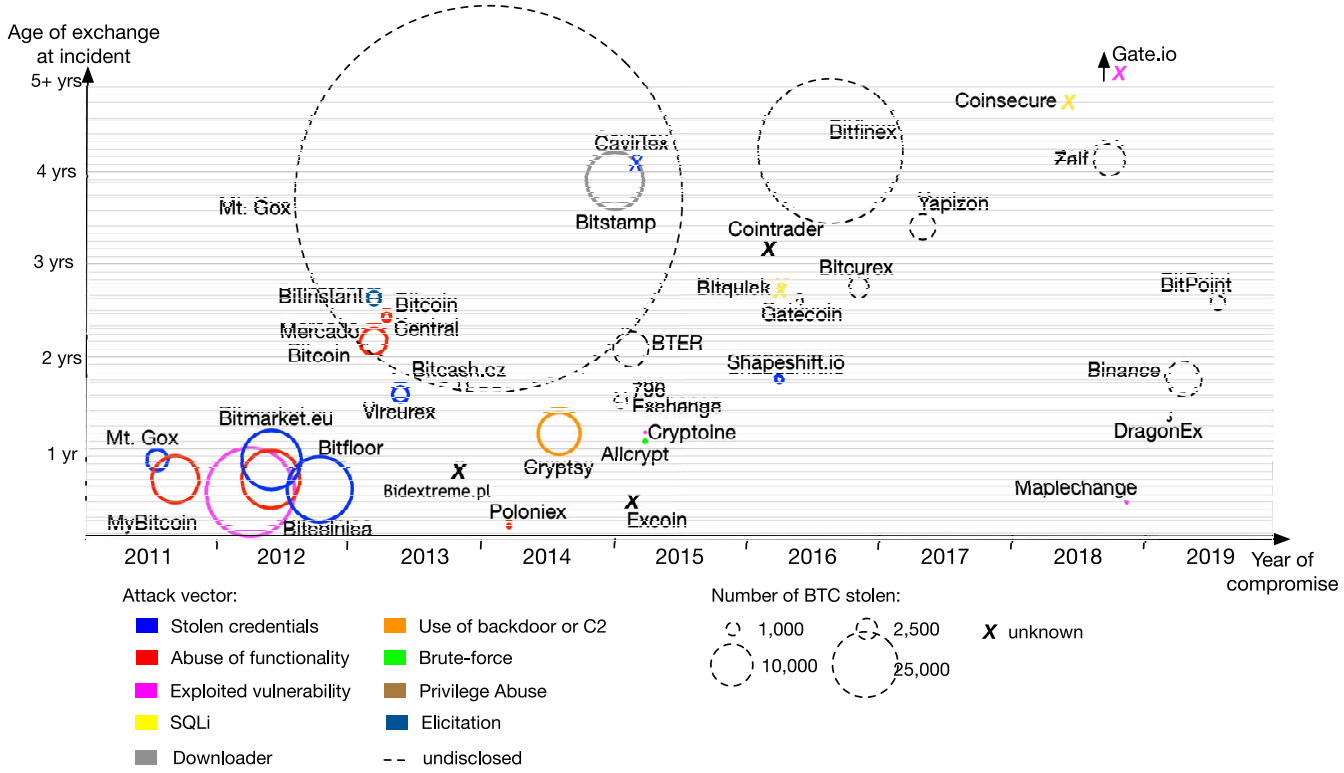

Fig. 1. Timeline of Bitcoin exchange breaches with respective attack vectors and BTC stolen.

from trivial exploitation of functionality or vulnerabilities to other hacking vectors.

Table I provides an overview of the breaches recorded in our dataset. The *Launch* and *Breach* columns denote when an exchange was first opened and breached subsequently, *BTC* how much BTC were lost, *USD* the loss in USD based on the average exchange rate in the breach month [22], *Action* and *Variety* how the attack vector used is classified within VERIS. In the *Attack* column, we have placed references to the sources used for the VERIS classification of each breach and further analysis in this section. The information in this column is based on reporting on forensic investigation by the breached party. When not available, we have drawn on reports from secondary sources such as media, emphasized in *italic*. *Closed* denotes whether an exchange closed as result of a breach. The asterisk indicates termination after a subsequent breach.

The most occurring varieties in our dataset of Bitcoin exchange breaches are: Unknown (12), Use of stolen credentials (6) and Abuse of functionality (5). In the sections below we will discuss the observations for these breach varieties in more detail, as well as our findings with regards to their impact.

A. Analysis of Attack Vectors

1) Increase of Unknown Variety Indicates Decrease in Disclosure of Breach Details: While we were able to identify the attack vector for most of the breaches in the first half of the time frame covered by our dataset, in recent years the communication of the TTPs used to breach exchanges has gotten more fuzzy. Of the 36 incidents recorded in our dataset, the specific attack vector remains unknown in 15 cases. Of these 15 cases, 9 appear in the last three years. In none of the latter cases, details on the vector through which the exchange was breached was not publicly disclosed.

While over recent years less exchanges terminate after a breach and theoretically better able to provide better incident response information, they often tend not to. While in the early years full details were usually not available due to the exchange going out of business (from 2011 until 2015, only 13.3% of breach methods remained undisclosed), over the last years exchanges survive but do seem to share less details out of concerns for their reputation. It is known that for various reasons, organizations are hesitant to share details on security breaches [23]. Other financial service organizations tend to exert more openness about breaches. One example of this is the hacking operation of Bangladesh Bank, which was disclosed officially by SWIFT [24], as well cyber security vendors [25]. This is partially explained by the fact that traditional financial institutions are subject to strict breach notification regulations such as the Federal Information Security Management Act of 2002 and the European General Data Protection Regulation. However, based on the information recorded in Table I, a trend of exchanges getting less transparent over the last couple of years can be observed.

2) Decrease in Use of Stolen Credentials Variety: The use of stolen credentials (*Stolen Creds* in VERIS terminology) is the method of choice in most of the early attacks in our dataset. In this type of attack the malicious actor breached the exchange platform and consequently exfiltrated funds by using privileged credentials. In many cases, credentials providing elevated (administrator) privileges were obtained through relatively low-level social engineering or

unsafely stored. While this type of breach does not involve exploitation of a vulnerability or another form of abuse, it is a cyber security breach because it affects system integrity.

Six exchanges in our dataset were breached through the use of stolen credentials, all of which occurred from 2011 up until 2016. In June 2011 Mt. Gox was breached with a compromised administrator account. More than 24,000 BTC were stolen from Bitfloor after the attacker managed to obtain credentials from their cloud provider to gain access to an unencrypted backup of a wallet used for cold storage. One explanation for the feasibility of this type attack is password reuse, because end users recycle the same password or variants of the same password through multiple online services.

After Unknown hacking, the use of stolen credentials is the biggest hacking vector in our dataset of Bitcoin exchange breaches (17%). The same goes for other financial services recorded in the VCDB, in which for 24 of a total 158 incidents this vector is employed (15%). Based on this data, both verticals are targeted and consequently exploited using the same methodology. However, where in traditional financial services stolen *user credentials* are mostly used to steal funds from individual users, the exchanges in our dataset were breached with *administrative credentials*, which provide instant access to funds of multiple users. The fact that the use of stolen credentials is decreasing over recent years indicates that exchanges have increased their security hygiene.

3) Decrease in Abuse of Functionality Variety: Abuse of functionality was the attack vector of choice in 5 breaches. Just like *Use of stolen creds*, this method does exploit a platform's access mechanisms. However, rather than the exploitation of a technical vulnerability, the attacker abuses legitimate platform functionality of the exchange platform or its hosting partner. Examples of these are the use of a flawed password recovery or discount modules, which can generally be avoided by thorough unit testing.

This vector was dominant among breaches of early movers in the exchange ecosystem, breached between 2011 and 2014. At the same time this observation is in accordance with a general trend observed in cyber attacks. Over the past years, attackers tend not to deploy malware or custom exploits. If not necessary to accomplish their objectives, they prefer to use functionality native to a target system. This way they are "living off the land" (LOTL). While LOTL attacks are employed both by low-level and sophisticated actors, their feasibility by the misuse of native features usually implies lax monitoring or security audits at the side of the victim. Our dataset records 5 data breaches through abuse of functionality, which is 14% of the total. In VCDB, this is recorded in just 2 incidents, which is only 1.29% of incidents recorded for the financial services industry.

4) Relatively Limited Deployment of Advanced Methods: As discussed in earlier sections, most exchanges in our dataset were breached using relatively straightforward attack vectors. Only three exchanges were compromised through the use of advanced techniques. Cryptsy was breached due to the exploitation of an intentionally placed backdoor in an open-source software dependency. After a malicious actor took over ownership of the development of Lucky7Coin, he was able to place an IRC backdoor into the wallet code base, allowing full and unlimited access to funds stored in the wallet [61]. Gate.io was breached due to a breach at Statcounter, which allowed attackers to place code in the visitor counting script used by Gate.io. Both were targeted attacks, as they were the only instance in which this specific vulnerability in the dependency was exploited. Furthermore the breach of Bitstamp in January 2015 involved multi-staged and targeted malware according to leaked post-mortem reporting. Apart from these cases, the exchanges in our dataset were hacked with relatively straightforward vectors. Especially given the considerable financial impact, this is a characteristic unique to Bitcoin exchanges. If the breach methods are not advanced, it implies the level of technical security is very low. And if security of an exchange platform is low, the company will not be able to keep up against the sophisticated nation state actors by which they are targeted, as mentioned in the introduction.

B. Analysis of Impact

1) Same Vectors, Different Outcomes: In the sections above we have found that the vectors used to breach Bitcoin exchanges are similar to those targeting financial institutions. The real world outcomes are however very different. Where attackers targeting Bitcoin exchanges are always motivated to exfiltrate funds, this is less the case in breaches of traditional financial institutions. According to the DBIR 2019, in 43% of incidents personally identifying information (PII) is exfiltrated and credentials in 38% of breaches [18]. According to the same report, theft of funds mostly happens to physical tampering attacks against automated teller machine (ATMs), which have declined over the last couple of years. Although the breach TTPs overlap, results and implications of breaches hugely differ between these types of organizations.

2) Hot Wallets Remain the Weak Spot: Exchanges use so-called hot and cold wallets like storage providers differentiate between hot and cold storage. Hot wallets provide liquidity to quickly facilitate transactions that characterize an exchange; depositing and withdrawal of funds to convert fiat to digital assets or exchange between BTC and ERC-20 tokens. The amount of assets stored in hot wallets should be sufficient to provide quick liquidity. The largest share of user funds held in custody should be held in cold wallets, which are stored offline or airgapped, meaning isolated from the regular local network and outside networks and ideally requiring physical access.

Except for 2 cases, in all breaches the funds exfiltrated by the hackers were stored in a hot wallet. Because these wallets are connected to the Internet, private keys can be obtained by breaching the server on which the wallet is stored. Only BTER [62] and Coinsecure [53] reportedly had their cold storage breached. Storing only as much funds as necessary in hot storage is considered good cyber security hygiene, as the offline nature of cold wallets makes them more difficult to breach. This potential financial impact of a breach is significantly decreased if the attacker is only able to compromise hot storage with a constrained amount of funds required to meet liquidity needs.

Our dataset shows that hot storage only provides security when implemented correctly. In the early years, exchanges went insolvent after a breach because they stored practically all funds in hot storage. Recently, regulatory frameworks imposing strict requirements on the custody of exchanges have been introduced in many jurisdictions. The Hong Kong Securities and Futures Commission is one of the first movers among global financial regulators with new regulation. As of November 2019, it requires Hong Kong-based platform operators to store 98% of assets in cold wallets and 2% in hot wallets. It also requires platforms to minimize the number of transactions from cold wallets and to insure funds to cover for a hack [63]. Furthermore, several leaders in the cryptocurrency ecosystem have established the Cryptocurrency Certification Consortium (C4), which has released the CryptoCurrency Security Standard [64]. This standard provides guidance to cryptocurrency companies to implement information security frameworks such as ISO 27001.

The above provides a representation of growing technical security maturity of Bitcoin exchange platforms. It does not require much technical sophistication to put an exchange platform online, as one can build on various widely-available open source components. However keeping such infrastructure secure takes significant resources and only a limited pool of people has the knowledge and experience of exchange security. The sophisticated threat actors targeting exchanges however investigate ample time to find a vulnerability that provides them with a foothold. And they only need one in order to further escalate their access level.

3) Decreased Exchange Closure Due to Breaches: Being breached was equal to insolvency and subsequently closing down in most of the first incidents recorded in our dataset of breached Bitcoin exchanges. However over recent years, this is not the case anymore. In most recent cases, exchanges resumed business after a period of ceased trading post-breach. This is only partially good news for owners of Bitcoin. In some cases stolen BTC were reimbursed on a 1:1 basis (Binance, BitPoint), but in other cases the exchange refunded a fixed percentage of BTC (Zaif, DragonEx). More controversial is the issuing of "IOU" ("I Owe You") tokens by Bitfinex and Yapizon, as these tokens serve as non-negotiable, informal measures of debt and thus are not redeemable for the actual value lost in BTC or USD.

In all cases, it is an improvement that customers are not left absolutely empty-handed after a breach. Our dataset includes exchanges paying out of pocket for this such as Binance's user asset fund [65], as well as selling the company in order to raise enough money [66]. Moore *et al.* [11] have argued that high-volume exchanges have a better chance to continue operations after a breach. This is interesting, as the bigger exchanges might have deeper pockets for security spending and thus to fend off cyber attacks. This would drive rational customers to large platforms, which indeed shows both in the trading volume (CoinGecko) and the amount of BTC held in custody [5]. This is however also contrary to the decentralized philosophy described by Satoshi Nakamoto in the Bitcoin whitepaper.

4) Amount of BTC Seized Per Breach Is Decreasing, Relative USD Yield Increasing: As it can be observed from Figure 1 and Table I, over the last few years the relative amount of BTC stolen as part of exchange security breaches tends to decrease. Seizures exceeding 10,000 BTC were not uncommon in the early years, however - except for some outliers - this has decreased in 2018 and 2019.

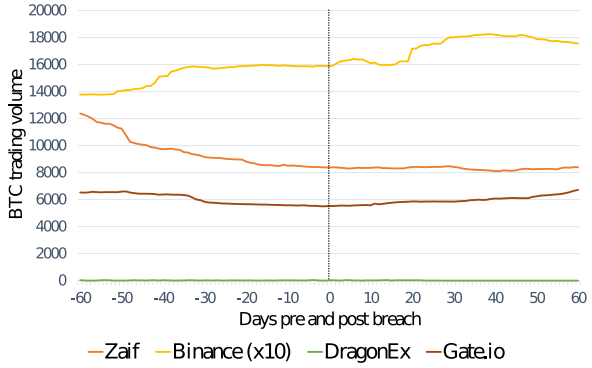

Fig. 2. Trading volume in BTC 60 days before and after breach.

Over recent years the relative USD yield has increased, while the BTC yield has decreased. This is a result of the increased BTC-USD exchange rate. Table I also shows that in the early years of our dataset, big amounts of BTC were taken through relatively low-level hacking vectors. Over recent years attack vectors have become more complex and less lucrative when simply considering BTC quantity. This could be the result of improved security best practices by exchange platforms, as well as improved incident response practices. Temporarily terminating withdrawal and trading functionality post-breach has become a common practice, as well as requesting other exchanges to which the hackers diverted BTC to freeze those.

5) No Impact on Trading Post-Breach: We have also found that in 4 recent breaches, the trading activity on the platform was not impacted. Figure 2 shows the reported trading volume 60 days before and 60 days after breach disclosure.

According to a report by Bitwise for the U.S. Securities and Exchange Commission [5], most Bitcoin exchange platforms operate wash trades or report fake inflated volume in order to increase attractiveness and exposure on market overview websites. This results in many cryptocurrency market overview websites reporting fake volume due to fabricated input data.

Historical trading volume data is not available for all exchange platforms, certainly not for those out of business. We have used data available from CoinGecko for exchanges which did not close after a breach as only then data is available. Although we had only BTC trade volume data available from 4 exchange platforms (Zaif, Gate.io, DragonEx, Binance), none of them did show a substantial impact on trading volume post-breach as can be seen in Figure 2.

DragonEx is at the bottom of the diagram, as its daily trade volume is below 100 BTC. The above shows that security breaches do not affect the trade volumes of Bitcoin exchange platforms that continue operations. We have to take into account that some amount of the volume might be fake. However the sustained trading can be explained by the incident response communication by these exchanges. They didn't

include much technical detail with regards to breach TTPs and every briefing assured customers that the exchange platform maintained control of the situation. This is coherent with the security industry best practice of frequent post-breach communication in order to avoid customer attrition [67].

VI. Evolution of Platform Security

The sheer fact that all organizations in our dataset were vulnerable to attacks in the past, invites the question whether platform security has improved over time. Data on this is scarce as many platforms have suspended their activities. Therefore we have decided to focus on 13 platforms from our dataset which are still operating. We have used scan results from Censys, which continuously scans the entire IPv4 address space since June 2016 for ports used in common services.

Adversaries usually deploy such port scans to acquire information about exposed services and the software versions behind them in a search for potential vulnerabilities for exploitation. To understand what services a Website was running at a given point in time, we used a historical passive DNS data source to resolve the domain names of these 13 platforms to the IP addresses used each day between June 2016 and December 2019, and correlated these with the banners grabbed from a particular IP address on that day. Additionally we have requested the HTTP Response Headers of each platform's *root www*. These are reconnaissance methods which do not negatively impact the platforms in question. To further comply with responsible disclosure of security issues, we do not identify individual platforms.

Based on the aspects observed, we see server security of the platforms investigated has significantly increased from 2016. As highlighted below, the Web security of these platforms has matured together with the overall exchange ecosystem, although bad security practices and security lapses are still comparatively widespread.

A. Multiple Services on Single IP

In our scans we have found multiple platforms running multiple services on the same IP address as the platform. This is not only bad practice due to availability and reliability, it also increases the attack surface as the potential pivot points for attackers are increased. If the mail relay is used to send spam emails, the Website ends up being blacklisted as well. Over the time frame from June 2016 until December 2019 this applied to 4 platforms in total, of which 1 still exposes FTP (which sends user credentials in plain text) and mail services at the moment of writing.

B. Vulnerable Web Server Versions

Given the fact that their Webpage is the critical frontend to exchange platforms and thus the single point of failure for such platforms, we were surprised to see that so many have been running vulnerable, unpatched versions of Web server software for long periods of time.

One of the bigger platforms has been running nginx 1.6.2 from at least June 2016 until February 2019, for which a high severity Denial of Service (DoS) vulnerability was reported already in February 2016 and a privilege escalation vulnerability in November 2016. Another high-volume platform was observed serving its platform's pages via Apache 2.2.22 from at least June 2016 until September 2017, which was then already vulnerable to several medium and high severity threats such as Cross-Site Scripting (XSS), buffer overflow, remote code execution, DoS and authentication bypass. In June 2019, one other platform had Microsoft Internet Information Server 8.5 exposed, for which a medium severity vulnerability was reported in 2014. Given a batch of only 13 platforms, observing 3 of them with major and easily exploited vulnerabilities post breach is astonishing and points to severe shortcomings in security practices such as vulnerability management.

C. Exposed Management Interfaces

The servers operating the exchange also need to be maintained and updated. These service interfaces would normally not be exposed to the public on the same IP as the Website itself, but rather be shielded in a separate compartment and accessible only to select sources. Already above we noted the existence of multiple services running on the platforms. From at least June 2016 until December 2019, one exchange had a Pure-FTPd service and Dovecot email server exposed. In September 2017, the same platform even exposed a vulnerable OpenSSH service version. In addition to configuration interfaces, another platform exposed a database server, MariaDB 5.5.5, a MySQL fork for 7 months over 2018 and 2019. In addition to normal architectural practices where databases would normally be placed deeper in the network and not accessible from a public network, this database was at that time already vulnerable to several high severity remote exploit vectors. Pure-FTPd was also exposed in 2016 and 2017 by another platform.

In principle, the attack surface could be minimized even when exposing such services to the outside, if mitigations such as IP address whitelisting can be applied. However residual attack surface still exists due to spoofing, and whitelisting was obviously not in place (or badly configured) if connections from an Internet scan service were permitted that would make the result available publicly. General security principles such as not to expose such services to untrusted zones like the Internet were thus not followed by the exchanges, even after breaches had happened.

D. Slow Patching in General

Based on our port scan data, it can be concluded that the deployment of security patches by exchange platforms has generally been slow. While over time less services are exposed in general, running vulnerable Web services is bad practice for online-only organizations. However room for improvement still exists, as later in this section we will show that as of now one platform is still running a vulnerable scripting engine.

E. HTTP Security Headers

When a Web server answers a request from a browser it includes HTTP response headers. A particular category of these are Security Headers, which instruct a browser how

to serve that particular page to avoid XSS and clickjacking. Implementing HTTP Security Headers is a straightforward and cost-effective measure, which protects against end users against malicious interventions. We have verified the configuration of the HTTP Security Headers by requesting the Webpage via *curl*, which is a non-intrusive method to better understand an organization's recognition of Web security. As shown in Table II, of the 13 platforms we investigated, we have found that three different headers were not properly configured for 6 out of the 13, and one type of header incorrectly for 7 out of the 13 platforms.

TABLE II
DETECTED IMPROPERLY CONFIGURED HEADERS

HTTP Security Header	Protects Against	Count
X-Content-Type-Options	Phishing, Cross-Site Scripting	6 of 13
X-Frame-Options	Clickjacking	7 of 13
X-XSS-Protection	Cross-Site Scripting	6 of 13
Strict-Transport-Security	Man-in-the-Middle	6 of 13

F. Server-Side Software

The *X-Powered-By* header is a non-standard HTTP response header, which can be manually adjusted to distract scanners targeting specific vulnerable software versions by increasing the version number to appear as more updated. Decreasing is not common, as it will only increase scrutiny.

For one platform, we found it running PHP/5.3.29, which dates back to 2015 and for which 5 critical vulnerabilities were disclosed in 2019. One of these vulnerabilities can be exploited to cause a buffer overflow, allowing for the server-side execution of attacker-controlled code. Given the history of attacks discussed earlier, this is a serious threat to the integrity of an exchange platform's transactions.

We have notified the platforms in question about open vulnerabilities. Based on our analysis there are however positive things to note. Over the time frame of the scans, we observed more platforms moving behind Cloudflare's reverse proxy service, which protects against several Web attacks. As of December 2019, all platforms except 2 have done this. Furthermore, less secondary services are getting exposed, although the practice still exists. While a causal relationship between security breaches which have hit the ecosystem so hard and the improved hygiene can obviously not be established, it seems that some maturing has happened in the overall ecosystem.

None of our findings are significant - any standard penetration test should uncover OWASP Top 10 vulnerabilities. We recommend exchange platforms not only to perform vulnerability scans and penetration tests on a regular basis, but to adhere to implement the OWASP Web Testing Guide in unit tests of their platform code. As simple technical vulnerabilities can threaten business continuity, over recent years many financial organizations have dedicated resources towards collecting threat intelligence on relevant attackers as part of their risk management and overall due diligence [68].

VII. TRACING STOLEN FUNDS

A question often asked after exchange breaches is how and to where BTC were funneled out. Such intelligence on post-breach TTPs provides an integral view of adversary behavior, however the analysis is delicate due to the privacy aspects inherent to Bitcoin's architecture. Bulk transaction analysis becomes even more complex due to the use of CoinJoin, third-party mixing services and the identification of services where funds terminate into fiat currency, which makes it difficult to establish a ground truth. The analysis is additionally complicated by the fact that details on laundering TTPs only become truly apparent when actors are indicted. In this section, we will analyze the post-breach TTPs to launder stolen funds from the Bitcoin exchanges in our dataset, based on official information shared by victim organizations and prosecutors.

A. Mt. Gox

The second breach of Mt. Gox in 2014 made the platform collapse and 850,000 BTC disappeared. According to U.S. justice department filings from 2017, the Russian-owned exchange platforms BTC-e and Tradehill were used to launder a significant portion of the Mt. Gox funds [69]. According to the Financial Crimes Enforcement Network, BTC-e was used to launder criminal money from miscellaneous origin, such as proceeds from Cryptolocker and Locky ransomware campaigns [70]. BTC-e failed to maintain effective AML measures and did not pursue any form of KYC, essentially favoring money launderers. According to the documents, the actors behind BTC-e allegedly managed to launder 4 billion USD. Among BTC-e's clients was the Russian state hacking group Fancy Bear [71], which is known to have used BTC to sponsor their hacking campaigns [72]. According to reporting by the BBC, most of BTC-e's user base vanished to a successor platform called Wex, of which all funds were allegedly seized by the Russian security service FSB [73].

The breach of Mt. Gox and the subsequent laundering has triggered several big investigations, some of which are still active. However the laundering TTP of using questionable exchanges to launder stolen funds, has become obsolete due to the restrictions raised by AML, KYC and FATF regulations.

B. Bitfinex Breach

With Bitfinex being a high-volume exchange, it would become the biggest breach after Mt. Gox. In June 2019 two Israeli individuals linked to the Bitfinex hack were arrested [74], but funds are still on the move. Many details on the adversary and its TTPs remain unclear. In August 2020 Bitfinex offered a 5% share of the assets recovered to anyone who puts the company in touch with the attackers and a 25% share to anyone who demonstrates control of the attacker wallets [75]. Just like with Mt. Gox, 4 years after the Bitfinex attack, the incident response is still ongoing.

A list of transactions and wallets associated with the hack was circulated by a Bitfinex director few days after the hack [76]. We have used this to analyze first and second-order movement of funds, in other words movement from wallets included in the list shared by the victim organization. As adversaries are known to use CoinJoin wallets and third-party

Fig. 3. Force-clustered transactions (green) from wallets (yellow) associated with Bitfinex attack from August 2016 until mid June 2020.

mixing services, tracing stolen funds on the blockchain is ambiguous and potentially unreliable. To keep analysis empirical, we focused on movement relating to addresses from the list distributed initially. In general, we have observed laundering the funds has been very laborious to the actors, with relatively limited success.

Figure 3 is a force-directed graph of outbound transactions from Bitfinex to the 410 wallets that have been active until now, based on the list shared by Bitfinex. A yellow dot represents a Bitcoin wallet, a green dot a mainnet transaction. The yellow dot in the center represents Bitfinex, surrounded by the wallets to which funds were exfiltrated as officially reported by Bitfinex. The main takeaway from the graph is how attackers diffuse funds to many wallets through a Web of transactions. As the nature and ownership of these wallets can only be speculated, the graph shows how the attackers use the Bitcoin network to obscure movement of yields from the breach. The list shared by Bitfinex contains a total of 2072 transactions associated with the hack, totalling 119.755 BTC. These transactions all took place between 8:54:54 and 12:18:35 on 2 August 2016. In total, we have recorded 1001 transactions associated with the Bitfinex hack, from August 2016 until June 2nd, 2020. Based on our analysis of these transactions, we have observed the following.

Manual obfuscation: In January 2017, the actors collected funds from several small wallets into a single 93 BTC wallet. This was then funneled to a wallet with another 15 BTC of stolen funds and then split into smaller quantities [77]. For the coming months, the funds then got separated into smaller quantities. While from that point onwards it cannot be established if the funds are still in possession of the attackers, apparent manual obfuscation of funds is a widely-documented laundering TTP [78].

Mixing: The attackers can be observed using supposed mixing wallets. One address has 394 incoming transactions from wallets mentioned in the official list [79]. This wallet was created on October 20, 2015 and already handling transactions prior to the Bitfinex hack. While the character of this wallet cannot be established, it has characteristics of a mixing wallet as it exclusively handled transactions worth few satoshis.

Recent activity: In a recent instance of consolidation, between June 1 and June 7, 2020, the attackers moved funds originating from several wallets with smaller holdings, also directly related to the hack [80]. A few days before, on April 28, 2020, the actors emptied a single wallet filled with 168 BTC directly after the attack, which is shown in the cluster just below the center of Figure 3.

The list shared by Bitfinex contains 2072 transactions associated with the hack, in total 119,755 BTC. The transactions took place between 8:54:54 and 12:18:35 on August 2, 2016. From this, 1001 subsequent transactions can be observed taking place from August 2016 until June 7, 2020. At the moment of writing, only 410 wallets of total 2072 have been depleted. The aggregate outbound transactions account for 2663 BTC, which is just 2.2% of the BTC stolen. While it is speculated what is causing this, it is evident that laundering stolen BTC - especially transferring Bitcoin to fiat assets - has become a complex operation for the adversaries behind the hack of Bitfinex.

TABLE III
TTPs IN LAUNDERING OF STOLEN BTC

Exchange	Breached	Laundering TTPs
Mt. Gox	2014-02	Fraudulent exchange
Bitfinex	2016-08	Mixing, manual mixing/splitting
Yapizon	2017-04	Bank withdrawals, money mules, gift cards

C. Yapizon

The North Korean state-sponsored Lazarus Group is associated with several cryptocurrency-related attacks [81]. The group is suspected of being behind the hack of Yapizon in 2017, recorded in our dataset. Furthermore the group is suspected of stealing altcoin from several other exchanges, as well as malware-based attacks to steal key pairs of unwitting users. In March of 2020, U.S. government authorities indicted two individuals associated with laundering of proceeds from these breaches [82]. From this recent case, it can be observed how adversaries launder stolen BTC in an increasingly regulated ecosystem. Compared to the laundering TTPs discussed earlier, laundering process has become more laborious, as the actors were observed withdrawing 34 million USD of stolen funds from a Chinese bank account that is linked to an exchange account. More interestingly, they were also observed converting 1.4 million USD worth of BTC into iTunes gift cards, which were used to purchase then-laundered BTC.

The practice of laundering through iTunes gift cards is striking, as gift cards are a known money laundering avenue. It is therefore astonishing such volumes have not raised any flags, or that the Chinese authorities did not intervene in the transfer or withdrawal. This shows that persistent actors will seek to maneuver around limitations put up in the ecosystem by bona fide actors.

According the analysis of these 3 cases, it is fair to establish that straightforward conversion of stolen BTC directly into fiat currency is a practice of a time gone by. Most exchanges have flagged wallet addresses associated with security breaches. Laundering criminally obtained BTC is further complicated by KYC and AML regulations. Over 4 years after the fact, few wallets associated with the Bitfinex hack have been emptied, with regulatory scrutiny only increasing.

VIII. Limitations

We have made significant effort to include all security breaches of Bitcoin exchanges in our dataset. Like any analysis driven by open source data, only publicly disclosed breaches can be included. Breaches may not be reported or remain unknown even to the victim. The composition of the dataset depends on the reporting obligation or generally responsible practice of breached parties. Despite this limitation we believe that our dataset is an accurate representation of Bitcoin exchange breaches over the eight years past. Analysis of an ecosystem as a whole provides a better reflection of reality than analysis of individual breaches. Our analysis serves as an analysis of its current state, as the trading of Bitcoin is an ecosystem in constant flux.

We have based our analysis on official reports, but in some cases strong rumors of exit scams exist. As indicated in the previous section, we have included these in Table I in the interest of completeness. Fear, uncertainty and doubt (FUD) are inherent to the Bitcoin community [83] and hence a tacit limitation to any research of the environment.

Furthermore, our classification of breaches is based on current representation of facts in official sources. The reporting of Bitcoin exchange breaches tends to be very light on technical detail, if any. Almost six years after the fact, is still not publicly known how the Mt. Gox breach, with the highest-ever amount of stolen BTC, could have taken place. With exchange regulations in development in several jurisdictions, as well as court cases on breaches currently ongoing, future work is necessary as more details become available.

IX. Conclusion

With the amount of fiat currency flowing into the Bitcoin market, exchange platforms are an attractive target for cyber criminal actors. We have analyzed 36 instances of cyber security breaches of Bitcoin exchange platforms, cumulatively accounting for at least 1,156,399 BTC stolen from their legitimate owners. Each of these incidents was facilitated by cyber security of the exchange platform and not the negligence of its users, the legitimate Bitcoin owners.

We have found that in recent years exchanges tend to disclose less technical details on the what and why of a breach compared to their earlier victims in our dataset. With regards to the vector used to breach exchanges, both the use of stolen credentials and the abuse of functionality are decreasing. This is good news, as a decrease of easy attack vectors suggest an increase in the levels of technical security of exchange platforms. Other positive developments are decreased exchange closure due to breaches and a decreasing amount of stolen BTC over recent years. This is partially due to other exchange platforms being willing to block funds directed to them by the attackers and returning those. Although the absolute BTC yield per breach has decreased, exchange platforms remain an interesting target due to increased BTC-USD exchange rate. Funds stored in hot wallets remain the primary target for attackers, as only 2 breaches in our dataset involved cold storage. The vectors used to breach Bitcoin exchanges overlap with those used in the broader financial services industry. Actual theft of funds is however rare in traditional financial services, where mostly personal information is targeted.

As of 2019, exchanges are required to comply with Know Your Customer and Anti-Money Laundering regulations in most jurisdictions. Compliance with such legislation is usually accompanied by stricter cyber security. However compared to other organizations in the financial sector, regulatory oversight on exchange platforms falls short in the protection of customer funds. A deposit insurance system, which provides customer protection in case an exchange becomes insolvent, can be a next step for the Bitcoin ecosystem. Exchange platforms can also take the lead, with recent cooperation in freezing and returning funds to breached exchanges serving as an example. The initiation of mutual aid agreements as prevalent in other industries can help formalize such arrangements.

Yet all cyber threats discussed in this article are a result of centralization of the ecosystem caused by centralized exchanges. Peer-to-peer trading is not vulnerable to these threats as decentralized exchanges do not store user assets. Decentralized trading however shifts security risk and thus responsibility to the user. Service providers keep building propositions on top of the Bitcoin technology stack, each with its own implications on the core attributes of confidentiality, integrity and availability. Whether risk is acceptable remains a responsibility of the user, who votes with his or her (Bitcoin) wallet.

References

[1] *Market Capitalization*, Blockchain.com, 2019. [Online]. Available: https://www.blockchain.com/charts/market-cap?timespan=2years

[2] *Weaknesses*, Bitcoin Wiki. [Online]. Available: https://en.bitcoin.it/wiki/Weaknesses

[3] S. Nakamoto. (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. [Online]. Available: https://bitcoin.org/bitcoin.pdf

[4] A. Antonopoulos. (2017). *Bitcoin Q&A: How Do I Secure My Bitcoin?*. [Online]. Available: https://www.youtube.com/watch?v=vt-zXEsJ61U

[5] "Analysis of real bitcoin trade volume," Rep., Bitwise, Schaumburg, IL, USA, 2019. [Online]. Available: https://static.bitwiseinvestments.com/Research/Bitwise-Asset-Management-Analysis-of-Real-Bitcoin-Trade-Volume.pdf

[6] "Report of the panel of experts established pursuant to resolution 1874 (2009)," UN Security Council, New York,, NY, USA, Rep. S/2019/171, 2019.

[7] K. Oosthoek and C. Doerr, "From hodl to heist: analysis of cyber security threats to bitcoin exchanges," in *Proc. IEEE Int. Conf. Blockchain Cryptocurrency*, 2020, pp. 1–9.

[8] M. Conti, K. E. Sandeep, C. Lal, and S. Ruj, "A survey on security and privacy issues of bitcoin," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 4, pp. 3416–3452, 4th Quart., 2018.

[9] I. K. Lim, Y. H. Kim, J. G. Lee, J. P. Lee, H. Nam-Gung, and J. K. Lee, *The Analysis and Countermeasures on Security Breach of Bitcoin* (Lecture Notes in Computer Science), pp. 720–732, 2014.

[10] A. Feder, N. Gandal, J. T. Hamrick, and T. Moore, "The impact of DDoS and other security shocks on Bitcoin currency exchanges: Evidence from Mt. Gox," *J. Cybersecurity*, vol. 3, no. 2, pp. 137–144, 2017.

[11] T. Moore, N. Christin, and J. Szurdi, "Revisiting the risks of bitcoin currency exchange closure," *ACM Trans. Internet Technol.*, vol. 8, no. 4, p. 50, 2018.

[12] *A Comprehensive List of Cryptocurrency Exchange Hacks*, Selfkey, Port Louis, Mauritius, 2019. [Online]. Available: https://selfkey.org/list-of-cryptocurrency-exchange-hacks/

[13] *A Huge List of Cryptocurrency Thefts*, Hackernoon, 2019. [Online]. Available: https://hackernoon.com/a-huge-list-of-cryptocurrency-thefts-16d6bf246389

[14] A. Shostack, "Experiences threat modeling at Microsoft," in *Proc. CEUR Workshop*, pp. 1–11, 2008.

[15] *Common Attack Pattern Enumeration and Classification (CAPEC)*, MITRE Corp., McLean, VA, USA. [Online]. Available: https://capec.mitre.org/

[16] K. Oosthoek and C. Doerr, "SoK: ATT & CK techniques and trends in windows malware," in *Proc. 15th EAI Int. Conf. (SecureComm)*, 2019. [Online]. Available: https://krisk.io/post/attack/

[17] *VERIS: The Vocabulary for Event Recording and Incident Sharing*, Verizon, New York, NY, USA, 2019. [Online]. Available: http://veriscommunity.net/

[18] *2019 Data Breach Investigations Report*, Verizon, New York, NY, USA, 2019.

[19] (2019). *The VERIS Community Database*. [Online]. Available: https://github.com/vz-risk/VCDB

[20] *VCDB JSON Directory README*, Verizon, New York, NY, USA. [Online]. Available: https://github.com/vz-risk/VCDB/tree/master/data/json

[21] *Top 100 Coins by Market Capitalization*, CoinGecko, 2019. [Online]. Available: https://www.coingecko.com/en

[22] *Bitcoin Historical Data*, Investing.com, Tokyo, Japan. [Online]. Available: https://www.investing.com/crypto/bitcoin/historical-data

[23] T. Ring, "A breach too far?" *Comput. Fraud Security*, vol. 2013, pp. 5–9, Jun. 2013.

[24] *SWIFT Customer Communication: Customer Security Issues*, SWIFT, Sydney, NSW, Australia, 2016. [Online]. Available: https://www.swift.com/insights/press-releases/swift-customer-communication_customer-security-issues

[25] "Lazarus under the hood," Rep., Kaspersky Lab, Moscow, Russia, 2018.

[26] M. Karpeles. (2011). *Clarification of Mt Gox Compromised Accounts and Major Bitcoin Sell-Off*. [Online]. Available: https://web.archive.org/web/20110919162635

[27] T. Williams. (2011). *MyBitcoin Incident Report—August 5th 2011*. [Online]. Available: http://archive.is/LUMzs

[28] *Manager Security Incident*, Linode, Philadelphia, PA, USA, 2012. [Online]. Available: http://archive.is/tRQ9

[29] *[Emergency ANN] Bitcoinica Site Is Taken Offline for Security Investigation*, Bitcointalk.org, 2012. [Online]. Available: https://bitcointalk.org/index.php?topic=81045.0

[30] *BitMarket.Eu Has Closed Down*, Bitcointalk.org, 2013. [Online]. Available: https://bitcointalk.org/index.php?topic=5441.msg1533170#msg1533170

[31] *Bitcoin Exchange BitFloor Suspends Operations After $250,000 Theft—The Verge*, The Verge, Washington, DC, USA, 2012. [Online]. Available: https://www.theverge.com/2012/9/5/3293375/bitfloor-bitcoin-exchange-suspended-theft

[32] *Events of Friday—BitInstant Back Online—Blog—Genesis Block—The BitInstant Blog*, BitInstant, New York, NY, USA, 2013.

[33] L. César. (2013). *Problema Do Mercado Bitcoin*. [Online]. Available: https://bitcointalk.org/index.php?topic=160150.0

[34] *Les Explications De Bitcoin Central*, Bitcoin-Central, 2013. [Online]. Available: https://bitcoin.fr/les-explications-de-bitcoin-central/

[35] "May 2013 report," Vircurex, Rep., 2013. [Online]. Available: https://vircurex.com/Reports/2013-05.pdf

[36] *Czech Bitcoin Exchange Bitcash.cz Hacked and up to 4,000 User Wallets Emptied*, CoinDesk, New York, NY, USA, 2013. [Online]. Available: https://www.coindesk.com/czech-bitcoin-exchange-bitcash-cz-hacked-4000-user-wallets-emptied

[37] *Polish Bitcoin Exchange Bidextreme.pl Hacked, Bitcoin and Litecoin Wallets Emptied*, CoinDesk, New York, NY, USA, 2013. [Online]. Available: https://www.coindesk.com/hacker-attack-polands-bitcoin-exchange

[38] *Japanese Bitcoin Heist "an Inside Job," Not Hackers Alone*, Daily Beast, New York, NY, USA, 2017. [Online]. Available: https://www.thedailybeast.com/japanese-bitcoin-heist-an-inside-job-not-hackers-alone

[39] Busoni. (2014). *BTC Stolen From Poloniex*. [Online]. Available: https://bitcointalk.org/index.php?topic=499580

[40] *Cryptsy Blog—Announcement*, Cryptsy, New York, NY, USA, 2014. [Online]. Available: http://blog.cryptsy.com/post/137323646202/announcement

[41] *Details of $ 5 Million Bitstamp Hack Revealed*, CoinDesk, New York, NY, USA, 2015. [Online]. Available: https://www.coindesk.com/unconfirmed-report-5-million-bitstamp-bitcoin-exchange

[42] *Chinese Exchange Gets "Goxed" for 1,000 Bitcoins*, Coin Telegr., New York, NY, USA, 2015. [Online]. Available: https://cointelegraph.com/news/chinese-exchange-suffers-1000-btc-loss-in-uncertain-service-compromise

[43] *Latest News: Update on Withdrawals*, Cavirtex, Calgary, AB, Canada, 2015. [Online]. Available: https://web.archive.org/web/20150220001845/https://www.cavirtex.com/news

[44] *Excoin—Announcement*, Excoin, Coconut Creek, FL, USA. [Online]. Available: https://web.archive.org/web/20150215200218/https://exco.in/

[45] *What happened, and What's Going on—AllCrypt Blog*, AllCrypt, 2015. [Online]. Available: https://archive.is/2UY7e

[46] *Bitcoin Exchange Cryptoine Hacked*, ZDNet, San Francisco, CA, USA, 2015. [Online]. Available: https://www.zdnet.com/article/bitcoin-exchange-cryptoine-hacked/

[47] *Bitcoin Exchange Cointrader Shuts Down After Alleged Hack*, CoinDesk, New York, NY, USA, 2016. [Online]. Available: https://www.coindesk.com/bitcoin-exchange-cointrader-shuts-down

[48] *Names, Phone Numbers, and Emails Leaked in BitQuick Exchange Hack—Bitcoin News*, Bitcoin.com, 2016. [Online]. Available: https://news.bitcoin.com/names-phone-numbers-emails-leaked-bitquick-exchange-hack/

[49] *A Timeline: ShapeShift Hacking Incident—ShapeShift*, Shapeshift.io, 2016. [Online]. Available: https://info.shapeshift.io/blog/2016/04/19/blog-2016-04-19-timeline-shapeshift-hacking-incident/

[50] *Gatecoin | Official Statement Regarding Gatecoin Hot Wallet Breach*, Gatecoin, Hong Kong, China, 2016. [Online]. Available: http://archive.is/rcnG4

[51] *Security Breach on Bitfinex*, Bitfinex, Hong Kong, China, 2016. [Online]. Available: http://archive.is/CnjAn

[52] *Korean Bitcoin Exchange Yapizon Confirms $5 mln Hack, All Customers To Pay With Balances*, Coin Telegraph, New York, NY, USA, 2017.

[53] *Bitcoins Worth Rs 20 Crore Stolen from Exchange in India's Biggest Crypto Theft*, Econ. Times India, New Delhi, India, 2018.

[54] *Report on Suspension of Deposit and Withdrawal of Virtual Currency and Our Response*, PR Times, Tokyo, Japan, 2018. [Online]. Available: https://prtimes.jp/main/html/rd/p/000000093.000012906.html

[55] *Crypto Exchange Zaif Hacked In $60 Million Bitcoin Theft*, CoinDesk, New York, NY, USA, 2018. [Online]. Available: https://www.coindesk.com/crypto-exchange-zaif-hacked-in-60-million-6000-bitcoin-theft

[56] *Minor Crypto Exchange Pulls Off Exit Scam, Steals All User Funds*, Bitcointalk.org, 2018. [Online]. Available: https://bitcointalk.org/index.php?topic=5059683.0

[57] *Gate.io Will Stop Using Statcounter for Traffic Stats*, Gate.io, 2018. [Online]. Available: https://www.gate.io/article/16665

[58] *Singapore-Based Crypto Exchange DragonEx Has Been Hacked*, CoinDesk, New York, NY, USA, 2019. [Online]. Available: https://www.coindesk.com/singapore-based-crypto-exchange-dragonex-has-been-hacked

[59] *Binance Security Breach Update*, Binance, Birkirkara, Malta, 2019. [Online]. Available: https://www.binance.com/en/support/articles/360028031711

[60] *Bitpoint Exchange Hacked for $32 Million in Cryptocurrency*, CoinDesk, New York, NY, USA, 2019.

[61] *Cryptsy Blog—Announcement*, Cryptsy, 2016. [Online]. Available: http://archive.is/FfECg

[62] *BTER Claims $1.75 Million in Bitcoin Stolen in Cold Wallet Hack*, CoinDesk, New York, NY, USA, 2015. [Online]. Available: https://www.coindesk.com/bter-bitcoin-stolen-cold-wallet-hack

[63] "Position paper on Regulation of virtual asset trading platforms," Hong Kong Securities Futures Commission (SFC), Hong Kong, China, Rep., 2019.

[64] *CryptoCurrency Security Standard (CCSS)*, C4, 2019. [Online]. Available: https://github.com/CryptoConsortium/CCSS

[65] *Binance Is SAFU: 7 Ways We Secure Your Assets 24/7*, Binance, Birkirkara, Malta, 2019. [Online]. Available: https://www.binance.com/en/blog/307883269744750592/Binance-Is-SAFU-7-Ways-We-Secure-Your-Assets-247
[66] *Zaif Crypto Exchange Reveals Takeover in New Hack Refund Plan*, CoinDesk, New York, NY, USA, 2018. [Online]. Available: https://www.coindesk.com/zaif-crypto-exchange-reveals-takeover-in-new-hack-refund-plan
[67] T. Caldwell, "The true cost of being hacked," *Comput. Fraud Security*, vol. 2014, pp. 8–13, Jun. 2014.
[68] K. Oosthoek and C. Doerr, "Cyber threat intelligence: A product without a process?" *Int. J. Intell. CounterIntell.*, to be published.
[69] *Russian National and Bitcoin Exchange Charged In 21-Count Indictment for Operating Alleged International Money Laundering Scheme and Allegedly Laundering Funds From Hack of Mt. Gox*, U.S. Attorneys Office Northern District of California, San Francisco, CA, USA, 2017.
[70] *FinCEN Fines BTC-e Virtual Currency Exchange $110 Million for Facilitating Ransomware, Dark Net Drug Sales*, FinCEN, Vienna, VA, USA, 2017. [Online]. Available: https://www.fincen.gov/news/news-releases/fincen-fines-btc-e-virtual-currency-exchange-110-million-facilitating-ransomware
[71] *How the DOJ Indictment of Russian Hackers Is Supported by Blockchain Analysis*, Elliptic, London, U.K., 2018. [Online]. Available: https://www.elliptic.co/our-thinking/doj-indictment-russian-hackers-blockchain-analysis
[72] *U.S. and Russia Spar Over Accused Crypto-Launderer*, Organized Crime and Corruption Reporting Project, 2019. [Online]. Available: https://www.occrp.org/en/investigations/us-and-russia-spar-over-accused-crypto-launderer
[73] *Bitcoins in the "FSB Fund": How $450 Million Disappeared From Wex Crypto Exchange*, BBC Russia, London, U.K., 2019. [Online]. Available: https://www.bbc.com/russian/features-50420738
[74] *Bitfinex Hackers Arrested After Three Years*, ZDNet, San Francisco, CA, USA, 2019. [Online]. Available: https://www.zdnet.com/article/bitfinex-hackers-arrested-after-three-years/
[75] *Up to US$400 Million Reward for Return of Stolen 2016 Bitcoin*, Bitfinex, Hong Kong, China, 2020. [Online]. Available: https://www.bitfinex.com/posts/494
[76] *Txid and Bitcoin Addresses Connected to the Bitfinex Theft*, Reddit user zanetackett, San Francisco, CA, USA, 2016. [Online]. Available: https://np.reddit.com/r/BitcoinMarkets/comments/4wizgv/txid_and_bitcoin_addresses_connected_to_the/
[77] *Transaction*, Blockchain.com, 2017. [Online]. Available: https://www.blockchain.com/btc/tx/2fb1ad2aceb70d2235e7092559447ec493c21a0bff01f793b8c0161d5f5e92c9
[78] J. Hu. (2018). *Generate and Download Thousands of Bitcoin Wallets in a Minute or Two*. [Online]. Available: https://medium.com/coinmonks/generate-and-download-thousands-of-bitcoin-wallets-in-a-minute-or-two-d42ce73d77d8
[79] *Address*, Blockchain.com, 2019. [Online]. Available: https://www.blockchain.com/btc/address/19cj6xavuXErZE9vyob9jsB4AhQRGNZ9z2
[80] *Transaction*, Blockchain.com, 2019. [Online]. Available: https://www.blockchain.com/btc/tx/357376f5188054b46cdcf21328d7cbfcef75eface94bcdd8e6e54d4edd7ad8f2
[81] *Group-IB: 14 Cyber Attacks on Crypto Exchanges Resulted in a Loss of $882 Million*, Group-IB, Singapore, 2018. [Online]. Available: https://www.group-ib.com/media/gib-crypto-summary/
[82] *Treasury Sanctions Individuals Laundering Cryptocurrency for Lazarus Group*, U.S. Department Of The Treasury, Washington, DC, USA.
[83] B. Craggs and A. Rashid, "Misplacing trust in bitcoin information sources," in *Proc. Misinf. Misbehav. Min. Web Workshop Held Conjunction (WSDM)* 2018.

Kris Oosthoek (Member, IEEE) received the M.Sc. degree from Erasmus University, Rotterdam, The Netherlands. He is currently pursuing the Ph.D. degree with the Delft University of Technology, The Netherlands. He is a Senior Cyber Threat Intelligence Analyst with the Dutch Government. He is a Researcher with the Cyber Threat Intelligence Lab, Delft University of Technology. His area of expertise includes cyber threat intelligence, network security, malware analysis, and security operations.

Christian Doerr (Member, IEEE) received a joint Ph.D. degree in computer science and cognitive science from the University of Colorado at Boulder, USA. He is a Professor of Cyber Security and Enterprise Security and the Director of the Cyber Threat Intelligence Lab, Hasso Plattner Institute, Potsdam, Germany. His research focuses on network security, cyber threat intelligence, and situational awareness.