

Software weaknesses in cryptocurrencies create unique challenges in responsible revelations.

BY RAINER BÖHME, LISA ECKEY, TYLER MOORE, NEHA NARULA, TIM RUFFING, AND AVIV ZOHAR

Responsible Vulnerability Disclosure in Cryptocurrencies

DESPITE THE FOCUS on operating in adversarial environments, cryptocurrencies have suffered a litany of security and privacy problems. Sometimes, these issues are resolved without much fanfare following a disclosure by the individual who found the hole. In other cases, they result in costly losses due to theft, exploits, unauthorized coin creation, and destruction. These experiences provide regular fodder for outrageous news headlines. In this article, we focus on the disclosure process itself, which presents unique challenges compared to other software projects.¹⁵ To illustrate, we examine some recent disclosures and discuss difficulties that have arisen.

While Bitcoin is the best known, more than 2,000 cryptocurrencies are in circulation, collectively valued at \$350 billion as of August 2020.⁶ Figure 1 conceptualizes the landscape as a stack. While the details differ, at the lowest level, each cryptocurrency

system is designed to achieve common security goals: transaction integrity and availability in a highly distributed system whose participants are incentivized to cooperate.³⁸ Users interact with the cryptocurrency system via software “wallets” that manage the cryptographic keys associated with the coins of the user. These wallets can reside on a local client machine or be managed by an online service provider. In these applications, authenticating users and maintaining confidentiality of cryptographic key material are the central security goals. Exchanges facilitate trade between cryptocurrencies and between cryptocurrencies and traditional forms of money. Wallets broadcast cryptocurrency transactions to a network of nodes, which then relay transactions to miners, who in turn validate and group them together into blocks that are appended to the blockchain.

Not all cryptocurrency applications revolve around payments. Some cryptocurrencies, most notably Ethereum, support “smart contracts” in which general-purpose code can be executed with integrity assurances and recorded on the distributed ledger. An explosion of token systems has appeared, in which particular functionality is expressed and run on top of a cryptocurrency.¹² Here, the promise is that business logic can be specified in the smart contract and confidently executed in a distributed fashion.

The emergence of a vibrant ecosystem of decentralized cryptocurrencies has prompted proposals that leverage the underlying technology to construct new central bank currency² and corpo-

» key insights

- **Cryptocurrency software is complex and vulnerabilities can be readily, and anonymously, monetized.**
- **Responsible vulnerability disclosure in cryptocurrencies is hard because decentralized systems, by design, give no single party authority to push code updates.**
- **This review of case studies informs recommendations for preventing catastrophic cryptocurrency failures.**



rate electronic money, such as Facebook’s asset-linked Libra. This article focuses on existing decentralized cryptocurrencies. Some lessons discussed here could also inform the design and operation of these prospective forms of digital money issued by public or private legal entities.

Bugs in cryptocurrencies. The cryptocurrency realm itself is a virtual “wild west,” giving rise to myriad protocols each facing a high risk of bugs. Projects rely on complex distributed systems with deep cryptographic tools, often adopting protocols from the research frontier that have not been widely vetted. They are developed by individuals with varying level of competence (from enthusiastic amateurs to credentialed experts), some of whom have not developed or managed production-quality software before. Fierce competition between projects and companies in this area spurs rapid development, which often pushes developers to skip important steps necessary to secure their codebase. Applications are complex as they require the interaction between multiple software components (for example, wallets, exchanges, mining pools). The high prevalence of bugs is exacerbated by them being so readily monetizable. With market capitalizations often measured in the billions of dollars, exploits that steal coins are simultaneously lucrative to cybercriminals and damaging to users and other stakeholders. Another dimension of importance in cryptocurrencies is the privacy of users, whose transaction data

is potentially viewable on shared ledgers in the blockchain systems on which they transact. Some cryptocurrencies employ advanced cryptographic techniques to protect user privacy, but their added complexity often introduces new flaws that threaten such protections.

Disclosures. Disclosures in cryptocurrencies have occurred in varying circumstances, from accidental discoveries, through analysis by expert developers and academics, to observing successful exploits in the wild. In the rest of this article, we highlight the difficulties and subtleties that arise in each case. The root causes of most of the difficulties lie in the special nature of cryptocurrencies: they are based on distributed systems that were designed to be difficult to change in order to provide strong guarantees on their future behavior. In order to change these rules, the consent of many participants is needed—participants who are often anonymous, and who are organized loosely in communities without governing bodies or regulatory oversight.

Here, we briefly highlight the differences between conventional software development and cryptocurrencies with regard to vulnerability disclosure, we identify key issues in the disclosure process for cryptocurrency systems, and we formulate recommendations and pose open questions.

How Is Disclosure Different?

Responsible vulnerability disclosure in cryptocurrencies differs from the con-

ventions adopted for general software products in several ways. Two fundamental differences arise from the very nature of cryptocurrencies.

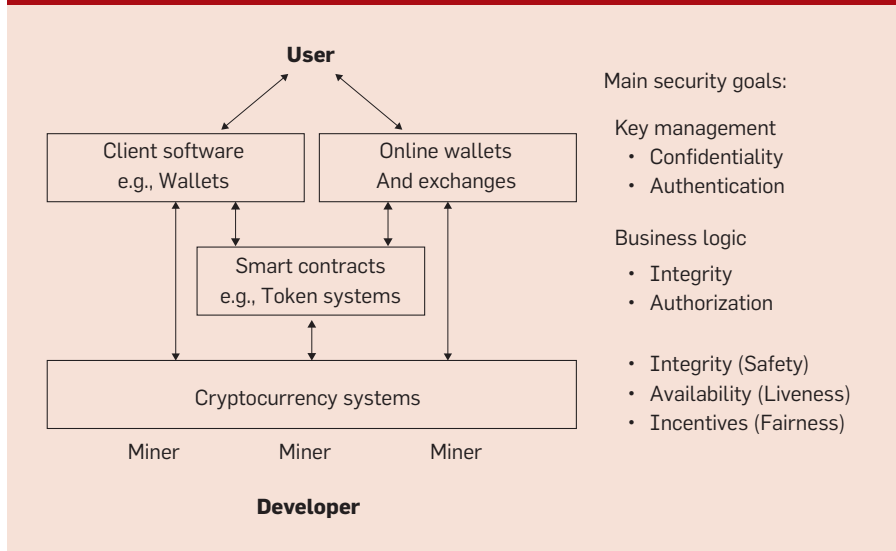
First, the decentralized nature of cryptocurrencies, which must continuously reach system-wide consensus on a single history of valid transactions, demands coordination among a large majority of the ecosystem. While an individual can unilaterally decide whether and how to apply patches to her client software, the safe activation of a patch that changes the rules for validating transactions requires the participation of a large majority of system clients. Absent coordination, users who apply patches risk having their transactions ignored by the unpatched majority.

Consequently, design decisions such as which protocol to implement or how to fix a vulnerability must get support from most stakeholders to take effect. Yet no developer or maintainer naturally holds the role of coordinating bug fixing, let alone commands the authority to roll out updates against the will of other participants. Instead, loosely defined groups of maintainers usually assume this role informally.

This coordination challenge is aggravated by the fact that unlike “creative” competition often observed in the open source community (for example, Emacs versus vi), competition between cryptocurrency projects is often hostile. Presumably, this can be explained by the direct and measurable connection to the supporters’ financial wealth and the often minor technical differences between coins. The latter is a result of widespread code reuse,²⁸ which puts disclosers into the delicate position of deciding which among many competing projects to inform responsibly. Due to the lack of formally defined roles and responsibilities, it is moreover often difficult to identify who to notify within each project. Furthermore, even once a disclosure is made, one cannot assume the receiving side will act responsibly: information about vulnerabilities has reportedly been used to attack competing projects,¹⁸ influence investors, and can even be used by maintainers against their own users.

The second fundamental difference emerges from the widespread design goal of “code is law,” that is, making code the final authority over the shared

Figure 1. Components of the cryptocurrency architecture covered in this article.



system state in order to avoid (presumably fallible) human intervention. To proponents, this approach should eliminate ambiguity about intention, but it inherently assumes bug-free code. When bugs are inevitably found, fixing them (or not) almost guarantees at least someone will be unhappy with the resolution. This is perhaps best exemplified by the controversy around the DAO, an Ethereum smart contract with a reentrance bug that was exploited to steal coins worth around \$50 million. After a community vote, the Ethereum developers rolled out a patch to reverse the heist, which (maybe surprisingly) turned out to be controversial. While the patch was accepted by large parts of the ecosystem, it was strongly opposed by a minority of Ethereum users arguing that it is a direct violation of the code-is-law principle, and the controversy ultimately led to a split of the Ethereum system into two distinct cryptocurrencies Ethereum and Ethereum Classic.¹ Moreover, situations may arise where it is impossible to fix a bug without losing system state, possibly resulting in the loss of users' account balances and consequently their coins. For example, if a weakness is discovered that allows anybody to efficiently compute private keys from data published on the blockchain,¹⁶ recovery becomes a race to move to new keys because the system can no longer tell authorized users and attackers apart. This is a particularly harmful consequence of building a system on cryptography without any safety net. The safer approach, taken by most commercial applications of cryptography but rejected in cryptocurrencies, places a third party in charge of resetting credentials or suspending the use of known weak credentials.

Ironically, these fundamental differences stem from design decisions intended to enhance security. Decentralization is prized for eliminating single points of control, which could turn out to be single points of failure. Giving code the final say is intended to preserve the integrity of operations. However, what may benefit security at design time becomes a significant liability after deployment once vulnerabilities are found.

Besides these fundamental differences, responsible disclosure for cryp-



The decentralized nature of cryptocurrencies, which must continuously reach system-wide consensus on a single history of valid transactions, demands coordination among a large majority of the ecosystem.



tocurrencies is characterized by specific features of the domain. The interpretation of system state as money, with many exchanges linking it mechanically to the conventional financial system, makes it easier and faster to monetize bugs than for conventional software, where vulnerability markets may exist but are known to be friction-prone.²³ Moreover, the cryptocurrency ecosystem reflects conflicting world-views, which prevent the establishment of basic norms of acceptable behavior. For example, invalidating ransomware payments via blacklisting has reignited the debate over censorship and the rule of law.²⁶

Finally, we note a difference in emphasis over certain aspects of disclosure. The conventional responsible disclosure discussion has focused on balancing users' interests in defensively patching versus national security interests of weaponizing vulnerabilities,^{25,31} without regard to whether the affected software is open or closed source. By contrast, open source software and code reuse are central to disclosure issues in cryptocurrencies, whereas balancing national and individual security considerations has so far not been widely discussed.

Throughout the rest of the article, we illustrate these differences with real cases before we derive recommendations and point to open problems.

Case Studies

We now review selected case studies of cryptocurrency vulnerability disclosures, highlighting aspects that teach us about the difficulties in response. We employ a multi-perspective method in selecting and researching these cases, ranging from the authors' direct experience as disclosers, interviews with developers and cryptocurrency designers, and through public reports. Interviews with open-ended questions were conducted by telephone, in-person or by email. Attribution is given unless the subject requested anonymity. The novelty and heterogeneity of the problem precluded a more systematic approach, though we hope that those informed by our findings can do so in future investigations. We investigate coins both small and large, because even the top coins have experienced severe bugs. While the software development pro-

cesses for prominent coins are more robust, the cases will show that all coins experience challenges to disclosure not seen in traditional software projects. Figure 2 presents a stylized timeline of the cases presented.

Cryptocurrency systems. Zcoin. We start with Zcoin, a relatively little known cryptocurrency that has suffered from repeated disclosures. Zcoin was the first to implement the Zerocoin protocol,²² which uses zero-knowledge proofs to enable untraceable transactions. In February 2017, an attacker exploited a typo in C++ code¹⁷ (using the equality operator ‘==’ instead of the assignment operator ‘=’) to generate 403,050 coins out of thin air. The new coins had a market value of \$750,000 and inflated the currency supply in cir-

ulation by 37%. In principle, such attacks can remain unnoticed due to the zero-knowledge veil, but the sheer number of coins created combined with the attacker’s impatience eventually led to its discovery. Within hours, the Zcoin team demanded that trading halt at big exchanges, published a blog post, and asked mining pools to suspend processing zero-knowledge transactions. A patch was released within a day, but the zero-knowledge feature remained disabled, thereby temporarily freezing all untraceable funds. This issue was resolved after four days when a “fork” altering the fundamental transaction validation rules was adopted by a majority of the miners. Even so, the attacker was able to abscond with the loot.

Later in 2017, a team of researchers including author Ruffing found another vulnerability in Zcoin that allowed an attacker to “burn” money in transit, that is, ensure no one, including the sender, recipient, and attacker, can further spend the coins.³⁰ Remarkably, the root cause of this vulnerability was an overlooked attack vector in the design and security analysis of the underlying Zerocoin protocol. While money burning does not serve the attacker directly, the attacker could profit indirectly, for example, by betting on falling prices of the affected cryptocurrency (short selling) and then publishing or exploiting the vulnerability. We have no evidence that such short-selling activity did indeed take place.

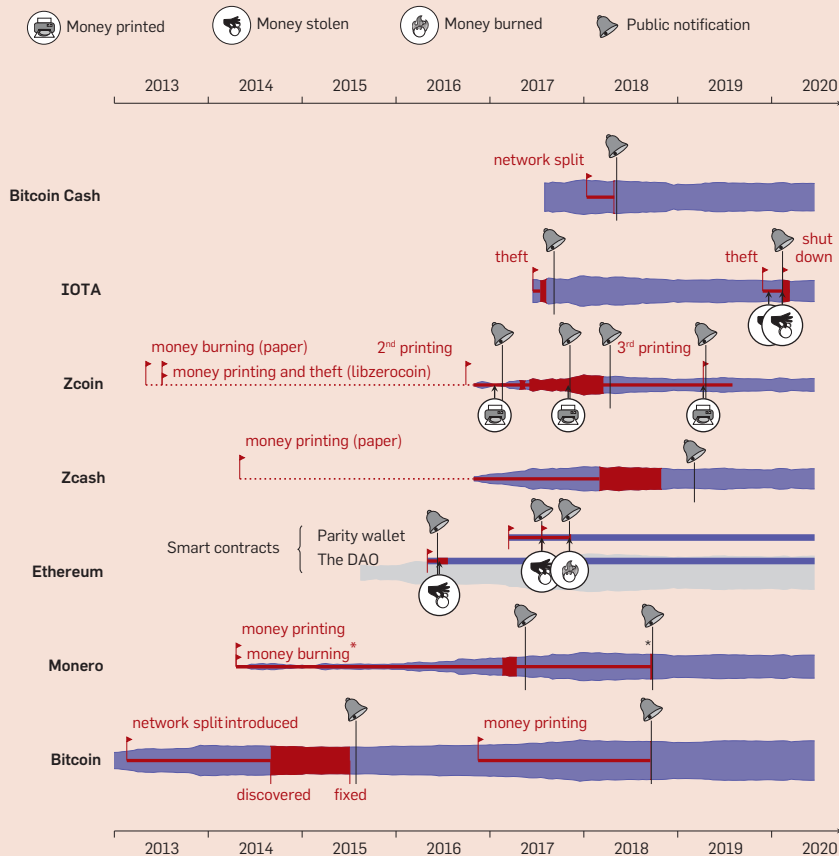
Having no cryptographer on its team, Zcoin hired Ruffing to provide advice and develop a patch. During the work, he identified two more vulnerabilities,²⁹ one enabling illegitimate coin generation and one allowing theft of money in transit. Both vulnerabilities stemmed from bugs in libzerocoin, a prototype library written by the inventors of the Zerocoin protocol for the purpose of validating their research. The Zcoin project had used that library as-is, despite the code’s prominent warning that the authors “are releasing this dev version for the community to examine, test and (probably) break” and that there are things that they have “inevitably done wrong.”²¹

Code reuse complicated the disclosure process of the three vulnerabilities.²⁹ Months after the initial notification, the discoverers found that more than 1,600 public GitHub repositories included verbatim copies of libzerocoin. Responsible and confidential disclosure to so many recipients is infeasible. Instead, the discoverers narrowed down the recipient set to less than 10 actual cryptocurrency projects, four of which they deemed trustworthy enough to be informed additionally. None of the projects had a clearly defined contact point or process for handling vulnerabilities.

Competition between projects prevented a coordinated response. For example, the notified project did not reveal to the reporters which of their competitors were also vulnerable. Coordination is essential because the first project to patch reveals the vulnerabili-

Figure 2. Visualization of the vulnerabilities discussed in this article.

The blue bars represent the underlying coins and their widths are proportional to their marketcap (for example, Coinmarketcap.org). The red bars visualize the discussed incidents from their introduction (flag) to their disclosure (wide bar) to their public announcement (bell). The additional symbol is used whenever money was stolen, burnt or printed.



ty, leaving the others unprotected. One currency was actually exploited in this way, and ironically, Zcoin itself was targeted because the patch was not adopted quickly enough. Dealing with the entire situation required tact and judgment by the discoverers, and the potential for every mistake to be catastrophic furthers the discoverers' burden.

As a result of the coin creation bugs, Zcoin improved continuous monitoring of aggregated balances, which led to the discovery of another creation bug in April 2019. The project repeated the notification process described earlier, disabled the zero-knowledge features via an emergency fork, and informed three potentially affected competitors. It took 10 days of investigation before a project developer identified the root cause in the design of the Zerocoin protocol. Unlike a simple implementation bug, there was no obvious way to fix the problem. The project's response was to migrate to an entirely different zero-knowledge protocol, suspending untraceable transactions in the meantime and freezing the affected funds until the new protocol was deployed in July 2019.³⁷

Zcash. Zcash, the commercial implementation of the Zerocash protocol,³ improves on Zerocoin's model for untraceable transactions. It too has suffered from similar issues.³² The proposal for the used algorithm for generating cryptographic material allowed a parameter to be published that should have remained secret. (Incidentally, a security proof was omitted because the scheme was similar to a previous one known to be secure.) The published value could have been used to undetectably generate coins out of thin air. The problem was discovered internally in March 2018 and fixed after 240 days in conjunction with a scheduled upgrade of the zero-knowledge protocol. Before and during the events the Zcash team had entered mutual disclosure agreements with the two largest competitors who reuse Zcash code. These competitors were notified two weeks after the fix with a schedule for public disclosure within a maximum of 90 days, which then took place in February 2019, almost one year after the discovery.³² Obscurity played a key role in this event: not only was the fix hidden in a larger update, the critical parameter was also



Unlike bugs in which coins are created, IOTA suffered a vulnerability that might have placed user funds at risk of theft.



removed from websites and a cover story spun around the "loss" of this piece of information. The intention of this obscurity was to protect Zcash's own interests and its users, as well as those of competing cryptocurrencies. On the downside, such long periods of obscurity may cast doubt on the trustworthiness of security claims in the future, and it remains unclear whether and to what extent the bug has been exploited.

Monero. The opposite of internal discovery is accidental public disclosure. This happened to Monero, the most popular implementation of the CryptoNote protocol.³⁵ In September 2018, an interested user posted a seemingly innocuous question to an online forum: "What happens if somebody uses a one-time account twice?" (paraphrased by the authors).⁷ Surprisingly, there was no protection against this action in the protocol. The revealed vulnerability allowed attackers to burn other people's funds. The problem was fixed within 10 days without known incidents and publicly announced thereafter.

A more serious vulnerability in the CryptoNote protocol affected all cryptocurrencies based on it. A post on a specialized cryptography mailing list in February 2017 revealed an issue, which implied a coin generation vulnerability in CryptoNote's basic cryptographic scheme.²⁰ The Monero team took note and developed a patch within three days and shared it privately with preferred parties, such as mining pools and exchanges. The true purpose of the patch was disguised in order to protect the rest of the users who were running vulnerable clients. After a fork to the validation rules that completely resolved the issue in Monero in April 2017, the Monero team informed other CryptoNote coins privately. One such coin, Bytecoin, was exploited immediately afterward, resulting in the illegitimate generation of 693 million coins.¹⁸ In a public disclosure that took place 15 days later, the Monero team described the aforementioned process and named unpatched competitors, including Bytecoin²⁰ (though Bytecoin claims that a patch had been issued to miners immediately after the exploit¹⁸). Perversely, the public disclosure attracted other investors to bid up the Bytecoin price. Its market capitalization grew

five-fold, briefly jumping into the top 10 cryptocurrencies by value. It remains unclear who exploited the bug, but Bitcoin holders certainly benefited from the price rise.

IOTA. Unlike bugs in which coins are created, IOTA suffered a vulnerability that might have placed user funds at risk of theft. Contrary to the best practice of using standardized cryptographic primitives, IOTA relied on a custom hash function that had a collision weakness.¹⁴ Author Narula and colleagues disclosed the vulnerability to the developers in July 2017. The vulnerability was patched by IOTA in August 2017 and made public by the disclosers in September 2017,¹³ offering several lessons about the disclosure process.

First, the vulnerability was fixed and deployed to the network quite quickly. On one hand, this is good because the potential vulnerability window is smaller. On the other hand, the speedy response was made possible due to the project's high level of control over the network, which runs contrary to the design goals of decentralized cryptocurrencies. Such control further allowed the operators to shut down its network to prevent theft from a vulnerable wallet for several weeks in early 2020.

The second lesson is that organizations may not respond favorably to a disclosure. Here, communications were tense, the existence and risk of the vulnerability was denied and downplayed, and the discoverers were threatened with lawsuits. The response echoes industry reactions to vulnerability disclosures related to digital rights management decades before.¹⁹ In the cryptocurrency case, there is a clear potential incentive conflict when the organization holds a large share of the coins and reasonably worries that the news could devalue holdings or prevent partnerships that might increase the value of holdings. Moreover, information about the bug could be exploited for profit by those possessing inside information about its existence prior to public disclosure.

Bitcoin Cash. Not to be confused with Bitcoin, "Bitcoin Cash" is derived from Bitcoin's codebase and was created due to disagreements within the ecosystem. Cory Fields, a contributor to the predominant implementation of Bitcoin, Bitcoin Core, was examining

change-logs of Bitcoin Cash's main implementation in April 2018.¹⁰ There he noticed that a sensitive piece of code dealing with transaction validation had been improperly refactored, causing a vulnerability. It would allow an attacker to split the Bitcoin Cash network, thereby compromising the consistency required for a cryptocurrency to operate.

As Fields noted, bugs like this cause systemic risk: if exploited, they could sink a cryptocurrency. The large amounts of money at risk prompt disclosers to take precautions. In this case, to protect his own safety, Fields chose to remain anonymous.¹⁰ The patching went smoothly, but we do not know if it would have been more contentious had he revealed his identity. Moreover, discoverers may want to demonstrate they behaved ethically, for example, that they sent a report to the developers. One possible mechanism is to encrypt the report with the developers' public key and publish the ciphertext and draw the developer's attention to it. This would require developers to provide public keys along with their security contact and have internal processes to handle incoming messages. Surprisingly, at the time Bitcoin Cash, a top-10 cryptocurrency worth billions of dollars, did not (though now they do). In our interview, Fields stressed he found it difficult to figure out what was the right thing to do. What helped him was to imagine the situation with swapped roles.

Bitcoin. A few months later, a developer from Bitcoin Cash disclosed a bug to Bitcoin (and other projects) anonymously. Prior to the Bitcoin Cash schism, an efficiency optimization in the Bitcoin codebase mistakenly dropped a necessary check. There were actually two issues: a denial-of-service bug and potential money creation.⁸ It was propagated into numerous cryptocurrencies and resided there for almost two years but was never exploited in Bitcoin.

This case teaches us three lessons: First, even the most watched cryptocurrencies are not exempt from critical bugs. Second, not all cases should be communicated to everyone in the network at the same time. The Bitcoin developers notified the miners controlling the majority of Bitcoin's hashrate of the denial-of-service bug first, making sure they had upgraded so that nei-

ther bug could be exploited before making the disclosure public on the bitcoin-dev mailing list. They did not notify anyone of the inflation bug until the network had been upgraded. Third, the disclosure involved deliberate deception of users: the Bitcoin developers published a patch describing it as only fixing the denial-of-service issue. This downplayed the severity of the bug, while at the same time motivating a prompt upgrade. This gave Bitcoin users and other affected cryptocurrencies time to adopt the fix, albeit with grumbling about the sudden public release. This highlights both a benefit and a downside to employing white lies in the disclosure process.

Silence is an alternative to white lies. The Bitcoin team took this option after an internal discovery in 2014. Bitcoin suffered from an inconsistency between different versions of the OpenSSL library. The 32-bit version was more tolerant in accepting variants of digital signatures than the 64-bit version, which could cause a loss of consistency if a signature is accepted only by the subset of nodes running on 32-bit. The mitigation turned into a year-long ordeal. Fixing OpenSSL was not an option, hence the stricter signature format had to be enforced in the Bitcoin codebase. Changes were made subtly and gradually in order to avoid drawing attention on the relevant piece of code. Users upgraded organically over a period of 10 months. The bug was made public when more than 95% of the miners had patched.³⁶

Smart contracts. Some cryptocurrencies, most prominently Ethereum, support "smart contracts." These are computer programs anyone can store on a shared blockchain, which then purports to guarantee correct execution. Contracts can receive, store, and send coins to users or other contracts according to their programmed logic. Smart contracts pose two further challenges to disclosure and patching. First, there is no club of miners whose incentives are aligned with the functioning of a specific contract. Therefore, relying on miners as allies to support smooth disclosure is usually not an option (though we will discuss an exception). Second, the code is not updatable by design to demonstrate commitment to the rules of operation,

hence the contract analogy. This may turn disastrous if the code contains bugs because machines, unlike arbitrators of real contracts, have no room for interpretation.

The DAO. The most famous example of a buggy contract is the DAO (short for Decentralized Autonomous Organization), the first code-controlled venture fund. Widely endorsed by an enthusiastic Ethereum community, in spring 2016 the DAO project collected user funds and stored them in a smart contract. Its visible balance of \$250 million (15% of all available coins at the time) made it a highly attractive target. It prompted scrutiny from security researchers who raised concerns,⁹ the closest activity to disclosure in the smart contract space we have seen. Three weeks later, an anonymous attacker managed to withdraw more than 3.5 million coins (about \$50 million) illegitimately from the DAO smart contract.¹ The attacker's trick involved making a small investment in the DAO, then withdrawing and thereby exploiting a re-entrance vulnerability in the refund mechanism. (The contract's bug was to not decrease the balance before sending coins, which in Ethereum passes control to the receiving party.) This exploit set off a vigorous debate over whether or not this behavior was abusive, since the code technically allowed the interaction.

The DAO incident could have been an example of an irreversible change of system state. However, the exceptional scale of the project and the involvement of the Ethereum community triggered a historic vote between miners to support a fork of the underlying cryptocurrency in order to "restore" the investments in the DAO contract. This intervention was highly controversial as it thwarted the very idea of immutable transactions, causing a group of purists to create a parallel instance, called Ethereum Classic, that was not rolled back. In hindsight, the incident raised the alarm to the smart contract community about the looming security issues. Today's contracts cannot hope for miner-enforced rollbacks because the uptake of the platform has diversified interests.

Parity wallet. Another example of a fund recovery, albeit partially successful, followed the Parity exploit in July

2017. The vulnerable contract implemented a multi-signature wallet, a mechanism that promises superior protection against theft compared to standard wallets. Intended uses include "corporate" accounts storing high value, such as the proceeds from initial coin offerings (ICOs). An anonymous attacker observed a discrepancy between the published and reviewed source code and the binary code, which was deployed for each of 573 wallets and omitted an essential access control step. This enabled a theft of coins worth \$30 million from three accounts. Parity discovered the attack as it was ongoing and published an alert. This would have enabled attentive users to rescue their funds (exploiting the same vulnerability) in a race against the attacker and imitators. At this point, a total of another \$150 million was essentially free to be picked up by anyone.³³ As expected, many users reacted slowly and found their funds missing. It turned out that a group of civic-minded individuals has taken the funds in custody in order to protect users and return them in a safe way. This example raises the question if protective appropriation of funds is legal or should even be expected from discoverers.

Users who nevertheless continued to trust the Parity wallet software were less lucky following a second incident. The Ethereum platform has a fuse mechanism that irrevocably disables code at a given address. In November 2017, a user (allegedly) inadvertently invoked this mechanism on a library referenced in 584 intentionally non-updatable contracts of the next-generation Parity wallet. A total of \$152 million was burned.³⁴ This time, no one intervened, presumably because the loss concerned only 0.5% of all coins.

We close by noting that as of this writing, we are not aware of any major

cases of responsible disclosures of vulnerabilities in smart contracts.

Recommendations and Open Questions

While best practices in secure software engineering and responsible disclosure¹⁵ are increasingly adopted in the cryptocurrency space, there always remains a residual risk of damaging vulnerabilities. Therefore, norms and eventually laws for responsible disclosure must emerge. What follows is a first step toward that end. Our synthesis of what can be learned from the cases is structured along three central issues of responsible disclosure: how to protect users, who to contact, when and how; and, how to reward the discoverer. The accompanying table sums up the recommendations outlined in this section.

How to protect users. *Discoverer safety.* If the vulnerability can make parties who may operate beyond the law substantially richer or poorer, the discoverer's personal safety should be considered.¹⁰ Death threats are not unheard of. Confidentially sharing the vulnerability with others the discoverer trusts (professional colleagues, notaries or the police) might reduce this risk. Sealed envelopes, or their digital variants such as time-locked encryption or secret sharing schemes, lessen the risk of unintended leakage. In addition, anonymous reporting may also reduce stress and tension. However, note that if the vulnerability is exploited, any proof the discloser knew of the vulnerability before its exploit could be used as evidence the discloser was the attacker.

Addressing vulnerable funds. If a vulnerability means that anyone can steal money from an account, should civic-minded defenders proactively steal to protect funds, like in the Parity wallet case? This touches on unresolved legal

Synthesis of recommendations.

Do's	Provide point of contact including public key Liaise with competitors who share code
Don'ts	Single out vulnerable competitors Bug bounties in your own coin
Depends	Use obscurity and white lies during disclosure Notify all affected projects unless there is conflict Built-in notification and feature "kill" switches
Need for action	Clarify right or obligation to preventively move vulnerable funds Establish clearinghouse and coordinator

questions. If “code is law” is the guiding principle, moving vulnerable funds must be legal. But courts are bound to real-world norms which differ across jurisdictions and circumstances. For example, in many places only law enforcement can legally expropriate property, including crypto coins. Elsewhere, disclosers could be obligated to intervene rather than stand by and allow a crime to take place. To give the discoverer legal certainty, it is essential to settle the basic question whether the discoverer could face legal consequences if she takes such precautions or break the law if she has the power and does not. If opting to leave the matter to law enforcement, other complications arise: Which law enforcement agency has jurisdiction and sufficient authority and is allowed to act? Do all law enforcement agencies possess the technical capability to intervene in time?

Preparing the system for disclosure. Given the inevitability of vulnerabilities, one strategy is to implement features in the cryptocurrency itself to automatically notify affected users of significant problems. In fact, Bitcoin used to have such an alert system, which enabled trusted actors to disseminate messages to all users and even suspend transactions. Such alert systems prompt difficult questions of their own, like who can be trusted with that authority in a decentralized system? Also, the alert system itself could become the target of attack, in much the same way that an Internet “kill switch” could create more security problems than it solves. Incidentally, Bitcoin itself abandoned the alert system over such concerns.⁴ A similar idea is to incorporate a mechanism to turn off particular features if significant vulnerabilities are later found. Dash utilizes such a system that lets the holder of a secret key turn features on and off at will.²⁷ PIVX supports a similar mechanism to disable zero-knowledge transactions, which proved useful during the Zerocoin disasters.

Despite the benefits such features bring, they contradict the design philosophy of decentralization and might expose the privileged party to law enforcement requests. Supposing a cryptocurrency could overcome these challenges and develop mechanisms for disseminating protective instructions,

the question of how to contact the trusted party who takes the precaution remains.

Who to contact, when, and how. *Provide clear points of contact.* Many cryptocurrencies are designed to avoid relying on privileged parties with substantial control. Yet this is in effect required to support responsible disclosure. It can be difficult to determine who is “in charge” (assuming anyone is) and who can fix the bug. Best practices recommend that developers provide clear points of contact for reporting security bugs, including long-term public keys.¹¹ Developers who reuse code are advised to publish alongside their own contact information that of the original code to aid the search for affected projects.

Identifying the responder. All communication by the discoverer should serve the end of fixing the bug. This means the discoverer must notify the party who is in the best position to solve the problem. For example, if the vulnerability affects the cryptocurrency’s core implementation, then the developers are the natural responders. There is a long history of bugs in exchanges,²⁴ in which case they would respond. It is important to note that once the responder has taken responsibility, the discoverer should adopt a “need-to-know” practice until the risk is mitigated. Sometimes the natural choice for responder is missing or untrustworthy. In this case, the discoverer can also serve as responder, or delegate the responsibility to a third party.

Responder communication with stakeholders. Given the decentralized nature of cryptocurrencies, the responder is usually not in a position to unilaterally act to fix the bug. Instead, the responder must seek stakeholders’ support. This means communicating the right messages at the right time. It could be dangerous to tell the full truth right away, so the message may justifiably include obfuscation or even white lies. Different stakeholders might require varying levels of detail at particular points in time. For bugs that require certain transactions to be mined for successful exploitation, the responder might encourage miners to upgrade first in order to deploy a fix as fast as possible. Exchanges can suspend trading in order to limit price shocks as bad news breaks, or aid in blocking the

transport of stolen funds. In other instances, wallet developers must be notified first, in order to deploy patches to their software. It is good practice to publish an advisory detailing the course of events and clarify any obfuscation or lies after the risk is mitigated. This transparency could mitigate the erosion of trust resulting from deception.

Coordination among multiple responders. As illustrated in the cases explored here, vulnerabilities often affect multiple projects. It is up to the discoverer to decide where to send the report. The reporter should be transparent about who has been informed. The discoverer can work with the responders to ensure that everyone affected has been notified. Coordination among responders is essential. Patches should be deployed as simultaneously as possible across affected projects, since the patching and publication of vulnerability information would leave others exposed if no precautions were taken. In some circumstances, the responders are competitors, and their attitudes toward one another range from suspicion to hostility.

Dealing with untrustworthy responders. While in the traditional security world, it is considered not only common courtesy but professionally and ethically required to inform other projects about vulnerabilities before disclosing their existence publicly. In the cryptocurrency world, one must adopt a more adversarial mindset. If the discoverer does not find a trustworthy responder, she can take on that responsibility. While one might not expect the discoverer to fix the bug, she could nonetheless take steps to protect users.

The situation is further complicated when multiple projects share a problem, and some are not trustworthy or are hostile toward each other. It is unreasonably burdensome for a discoverer to adjudicate such conflicts. Responders can make a best effort to identify affected parties (for example, searching for coins sharing common codebases) and notify accordingly. This points to the need for developing a clearinghouse, à la CERT/CC.


External authorities. Banks, payment processors, and other key financial institutions are often required to report vulnerabilities to banking regulators, who can coordinate the response if

needed. There is no current equivalent for cryptocurrencies, and it is unclear under which jurisdiction such a thing would reside. Should some global reporting agency of this nature be formed? If so, how might it successfully operate given a community whose common ground is removing the need for central parties? An external body modeled on CERT/CC might serve as a useful starting point. A less formal and more decentralized example to consider is iamthecalvary.org, an initiative bringing together security researchers with medical device manufacturers to promote responsible vulnerability disclosure and remediation.

How to reward the discoverer. The article has shown that disclosing a cryptocurrency vulnerability and reacting responsibly is very burdensome. Interviewees have reported sleepless nights and fears for their safety, which in turn has altered their professional collaborations and friendships. The alternative to profit from the vulnerability, potentially anonymously, is tempting. This is why cryptocurrencies specifically cannot expect altruistic behavior and must instead incentivize responsible disclosure.¹¹

Bug bounties offer an established way to reward those who find bugs.⁵ It stands to reason they would be a natural fit for cryptocurrencies, given they have a built-in payment mechanism. However, denominating the reward in its own currency is problematic, since its value might diminish as a result of disclosing the vulnerability, and you are effectively rewarding the discloser in a currency which she just found to be buggy. Other approaches are possible—for example, Augur (a smart contract market platform) is experimenting with exploit derivatives. It is not unreasonable to think that the cryptocurrency community might innovate a solution that could be a model for the broader software community. Nevertheless, monetary rewards must complement and cannot substitute for healthy norms and a culture that welcomes vulnerability disclosure.

Acknowledgments. This article is a result of a breakout group at the Dagstuhl Seminar 18461, “Blockchain Security at Scale.” The authors thank C. Fields, M. Fröwis, P. van Oorschot, and their interview partners.

This work is partially funded by: Archimedes Privatstiftung, Innsbruck, U.S. National Science Foundation Award No. ~ 1714291, ISF grant 1504/17, HUJI Cyber Security Research Center Grant, DFG grants FA 1320/1-1 (Emmy Noether Program) and SFB 1119—236615297 (Crossing), and BMBF grant 16KIS0902 (iBlockchain). 

References

- Atzei, N., Bartoletti, M. and Cimoli, T. A survey of attacks on Ethereum smart contracts. In *Proceedings of the Principles of Security and Trust*. M. Maffei and M. Ryan, Eds. LNCS 10204 (2017), Springer, 164–186.
- Bech, M. and Garratt, R. Central bank cryptocurrencies. *BIS Quarterly Rev.* 9 (2017), 55–70.
- Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I. and Virza, M. Zerocash: Decentralized anonymous payments from Bitcoin. In *Proceedings of the IEEE Symp. on Security and Privacy*, 2014.
- Bishop, B. Alert key disclosure. Bitcoin development mailing list; <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2018-July/016189.html>.
- Böhme, R. A comparison of market approaches to software vulnerability disclosure. *Emerging Trends in Information and Communication Security*. G. Müller, Ed. LNCS 3995 (2006). Springer, Berlin Heidelberg, 298–311.
- CoinMarketCap. Global charts, 2019; <https://coinmarketcap.com/charts/>.
- dEBRUYN. A post mortem of the Burning Bug, 2018; <https://web.getmonero.org/2018/09/25/a-post-mortem-of-the-burning-bug.html>.
- Bitcoin Core Developers. CVE-2018-17144 Full Disclosure; <https://bitcoincore.org/en/2018/09/20/notice/>.
- Dino, M., Zamfir, V. and Sirer, G.E. A call for a temporary moratorium on the DAO. Blog post; <http://hackingdistributed.com/2016/05/27/dao-call-for-moratorium/>.
- Fields, C. Responsible disclosure in the era of cryptocurrencies. Blog post, 2018; <https://bit.ly/3cgSdYs>.
- Fields, C. and Narula, N. Reducing the risk of catastrophic cryptocurrency bugs. Blog post, 2018; <https://bit.ly/2SKbd9Y>.
- Fröwis, M., Fuchs, A. and Böhme, R. Detecting token systems on Ethereum. In *Proceedings of the Financial Cryptography and Data Security*. I. Goldberg and T. Moore, Eds. 2019.
- Heilman, E., Narula, N., Dryja, T. and Virza, M. IOTA vulnerability report: Cryptanalysis of the curl hash function enabling practical signature forgery attacks on the IOTA cryptocurrency, 2017; <https://github.com/mit-dci/tangled-curl/blob/master/vuln-iota.md>.
- Heilman, E., Narula, N., Tanzer, G., Lovejoy, J., Colavita, M., Virza, M. and Dryja, T. Cryptanalysis of curl-p and other attacks on the IOTA cryptocurrency. IACR Cryptology ePrint archive report 2019/344; <https://eprint.iacr.org/2019/344>.
- Householder, A., Wassermann, G., Manion, A. and King, C. The CERT Guide to Coordinated Vulnerability Disclosure. Special Report CMU/SEI-2017-SR-022.
- Hutchinson, L. All Android-created Bitcoin wallets vulnerable to theft. *Ars Technica*, 2018; <https://bit.ly/2SMHiy5/>.
- Insom, P. Zcoin's Zerocoin bug explained in detail. Blog post, 2017; <https://zcoin.io/zcoins-zerocoin-bug-explained-in-detail/>.
- Juarez, A.M. Fraudulent transactions allowed by the CryptoNote key image bug remain valid. Archived version of Bytecoin GitHub; <https://bit.ly/2WbYkrn>.
- Lok, C. Dispute over digital music muzzles academic. *Nature* 411, 6833 (2001), 5.
- Luigi1111 and Riccardo “fluffypony” Spagni. Disclosure of a major bug in CryptoNote based currencies. Blog post, 2017; <https://bit.ly/2xHiTmb>.
- Miers, I. README of libzerocoin, 2013; <https://bit.ly/2L94ORJ>.
- Miers, I., Garman, C., Green, M. and Rubin, A. Zerocoin: Anonymous distributed e-cash from Bitcoin. In *Proceedings of the 2013 IEEE Symp. on Security and Privacy*.
- Miller, D. The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales. In *Proceedings of the Workshop on the Economics of Information Security*. Carnegie Mellon University, Pittsburgh, PA, 2007.
- Moore, T., Christin, N. and Szurdi, J. Revisiting the risks of Bitcoin currency exchange closure. *ACM Trans. Internet Technology* 18, 4 (2018), 50:1–50:18.
- Moore, T., Friedman, A. and Procaccia, A. Would a ‘cyber warrior’ protect us: Exploring trade-offs between attack and defense of information systems. In *Proceedings of the New Security Paradigms Workshop*. A.D. Keromytis, S. Peisert, R. Ford and C. Gates, Eds. ACM, 2010, 85–94; <https://tylermoore.utulsa.edu/nspw10.pdf>.
- Möser, M. and Narayanan, A. Effective cryptocurrency regulation through blacklisting, 2019; <https://maltemoeser.de/paper/blacklisting-regulation.pdf>.
- Multi-Phased Fork. Glossary item in developer documentation. Dash project, 2017; <https://dash-docs.github.io/en/glossary/spork>.
- Reibel, P., Yousaf, H. and Meiklejohn, S. An exploration of code diversity in the cryptocurrency landscape. In *Proceedings of the 2019 Financial Cryptography and Data Security Conf.* I. Goldberg and T. Moore, eds.
- Ruffing, T., Thyagarajan, S., Ronge, V. and Schröder, D. A cryptographic flaw in Zerocoin (and two critical coding issues). Blog post, 2018; <https://bit.ly/2WAib2B>.
- Ruffing, T., Thyagarajan, S., Ronge, V. and Schröder, D. (Short Paper) Burning Zerocoins for fun and for profit—A cryptographic denial-of-spending attack on the Zerocoin protocol. In *Proceedings of the Crypto Valley Conf. on Blockchain Technology*, (Zug, Switzerland, June 20–22, 2018). IEEE, 116–119; <https://doi.org/10.1109/CVCBT.2018.00023>.
- Schwartz, A. and Knake, R. Government's Role in Vulnerability Disclosure. Harvard Kennedy School discussion paper, 2016.
- Swihart, J., Winston, B., and Bowe, S. Zcash counterfeiting vulnerability successfully remediated. Blog post, 2019; <https://bit.ly/2YD0790/>.
- Parity Technologies. The multi-sig hack: A postmortem, 2017; <https://www.parity.io/the-multi-sig-hack-a-postmortem/>.
- Parity Technologies. A postmortem on the parity multi-sig library self-destruct, 2017; <https://www.parity.io/a-postmortem-on-the-parity-multi-sig-library-self-destruct/>.
- van Saberhagen, N. CryptoNote v 2.0. White Paper, 2013; <https://cryptonote.org/whitepaper.pdf>.
- Wuille, P. Disclosure: Consensus bug indirectly solved by BIP66. Bitcoin development mailing list, 2015; <https://bit.ly/2zeC5rX>.
- Yap, R. Further disclosure on Zerocoin vulnerability. Blog post, 2019; <https://zcoin.io/further-disclosure-on-zerocoin-vulnerability/>.
- Zohar, A. Bitcoin: Under the hood. *Commun. ACM* 58, 9 (Sept. 2015), 104–113.

Rainer Böhme (rainer.boehme@uibk.ac.at) is a professor of computer science at Universität Innsbruck, Austria.

Lisa Eckey (lisa.eckey@tu-darmstadt.de) is a cryptographer at TU Darmstadt, Germany.

Tyler Moore (tyler-moore@utulsa.edu) is the Tandy Professor of Cyber Security and Information Assurance at The University of Tulsa, OK, USA.

Neha Narula (narula@mit.edu) is Director, Digital Currency Initiative, MIT, Cambridge, MA, USA.

Tim Ruffing (crypto@timruffing.de) is a cryptographer at Blockstream.

Aviv Zohar (avivz@cs.huji.ac.il) is an associate professor of computer science at Hebrew University Jerusalem, Israel.

Copyright held by authors/owners.
Publications rights licensed to ACM.



Watch the authors discuss this work in the exclusive *Communications* video. <https://cacm.acm.org/videos/vulnerability-disclosure>