securosys

White Paper

SECUROSYS

Safeguarding of Crypto Assets

Address

Förrlibuckstrasse 70
8005 Zurich
Switzerland

Phone & Fax

Phone: + 41 44 552 31 00

Fax: + 41 44 552 31 99

Email

Email: info@securosys.com

Website: www.securosys.com

Table of Content
Whitepaper

Summary 03

Introduction 04

Attack Vectors 04

Physical Access 04

Side-channel Attacks 06

Randomization Weaknesses 07

Quantum Computing 08

Unauthorized Operations 09

Methods for Failure Protection 12

Considerations: Regulatory and market 13

Trade-offs 14

Trust 14

Price 14

User (Developer) Experience 15

Summary

In this paper, we address various aspects of safeguarding cryptocurrencies and other crypto-assets. The focus is on security, scaling, insurability, and regulatory aspects.

We take a look at different technologies and methods available for custodial platform, ranging from open-source cold-storage standards using on-chain multi-signature schemes, through Secure Multi-Party Computation to Hardware Security Modules by both legacy manufacturer and Securosys.
We also explore advantages and shortcomings of Hardware Security Modules, how they must evolve to keep up with the paradigm change introduced by cryptocurrencies to help us understand how to lead that change.

This paper is primarily a result of our own research to identify whether our core business is still relevant to this rapidly growing industry to help us decide on our own strategy and a potential pivot in thereof, but we believe it is worth sharing the insights we gained to contribute to the important and ongoing debate and to collect feedback on the conclusions derived.
The unique part of the job when it comes to protecting cryptocurrencies as opposed to other PKI applications is the importance of protection against unauthorized use of otherwise physically inaccessible key material.

While a traditional multi-signature approach combined with hardware cryptocurrency wallets or with open-source cold-storage approaches such as the Glacier Protocol provides a significant improvement in the crypto asset protection for individuals, it is inefficient, inflexible and unscalable for enterprises. Multi-party computation addresses some shortcomings of both multi-signature and legacy Hardware Security Modules but remains inflexible in the face of increasingly stricter and more sophisticated security policy requirements mainly due to lack of time-specific approval rules.

We hence are confident that our blockchain-focused HSM will emerge as the leading solution in terms of security, scalability, flexibility, and developer experience, mainly thanks to its highly customizable transaction control mechanisms and its physical properties built on decades of industry's best practice.

Introduction

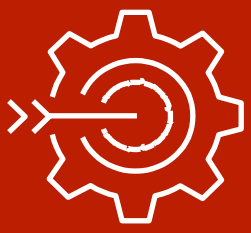

Rather than arbitrarily comparing solutions, we looked at various attack vectors and failure modes, and analyses how different approaches help custodians defend against and prevent such events. The types of vulnerabilities we analyzed are:

- Physical access to the key material;
- Unauthorized operations with the keys allowing the attacker to sign asset withdrawal transactions;
- Randomization algorithm weakness, which could be exploited to compute what should be an unguessable key;
- Calculation of private key data from its respective public key (something currently prevented by design of the asymmetric cryptography, but potentially feasible with future quantum computing);
- Hardware failure leading to a loss of private key material.

Attack Vectors
Physical Access

It is sufficient to simply copy the data that is the private key in order to sign a transaction transferring cryptocurrency balance. This can happen by a malicious external attacker gaining privileged remote access to the filesystem, by a hosting center administrator (ab-)using their physical access to the data storage, or by an employee copying the key to their private storage.

An offline approach with well-established operational procedures such as the Glacier Protocol provides sufficient protection for an individual long-term holder against such an attack. The user can store the key in a secure vault, encrypt it using a passphrase, use multi-signature and split the keys to multiple protected locations, or combine of all of these three methods. However, this is not possible where the business application is involved – the application must have the ability to conduct transactions semi-autonomously and therefore, the key must be accessible online or at least on a reasonably short notice. For real-time transfers, the key cannot be protected by a human passphrase input either.

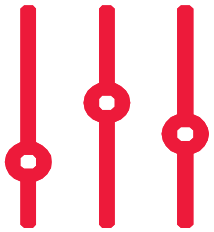

To allow for secure online availability of the asset balance while limiting physical exposure of the key material, the asset transactions can be made subject to multi-signature authorization where the signature authority can be split amongst multiple physically segregated systems. However, unless the physical storage of these systems is truly physically tamper-protected, the complexity of the attack grows merely in linear proportion to the size of the required quorum. The physical security could be theoretically strengthened by hardware key devices such if those were kept safe (since they are not sufficiently tamper-protected either) but given that these are not built to be highly reliable, their required redundancy introduces more options for their physical misuse.

Some solutions built on Secure Multi-Party Computation (SMPC) improve on the multi-signature approach by randomly shuffling the split key material frequently, which would require the attacker to gain access to all parts of the key at the same time. However, given that these solutions are software-based and have to be online, relying on unprotected forms of storage, the successful attack would require merely linearly proportional amount of work and patience on the side of the attacker.

It is indisputable, that despite their shortcomings in other areas, this one is where Hardware Security Modules stand out most since they build on decades of the best practices in preventing attackers from getting physical access to the data they protect. They achieve that using the following methods:

- Preventing a business application from copying the key material itself, allowing it to merely request operations with the keys - the operations are then conducted within the device and only the result of the operation is returned.

- Tamper-proof physical protection for the key storage using a set of sensors, which detect any attempt at gaining unauthorized access to the keys and in such a case render the storage unreadable.

Attack Vectors
Side-channel Attacks

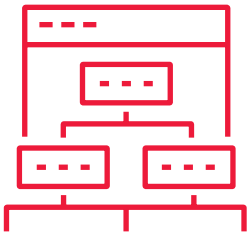

Side-channel attacks are alive and well

The recent disclosure of Plundervolt was just another example that side-channel attacks are alive and well, even on supposedly secure enclaves. These software vulnerabilities allow attackers to use manipulation or observation of physical properties of a hardware to extract what should be securely stored data, including private keys.

Multi-signature makes these types of attacks more complicated in an obvious way - it would be required to successfully run it on multiple machines of potentially different types in presumably different locations. Having at least one of the authorization devices offline in a physically segregated location minimizes the risk to virtually zero, but it comes at scaling and transaction speed costs.

The same goes for SMPC solutions with the additional advantage of some of its implementations, that the shared key material is being randomly rotated regularly and thus highly coordinated effort would be required for such attack to succeed.

HSMs are obviously purposefully built to prevent such attacks, especially Primus HSM by Securosys. Together with our research partner HSR, a multi-year research effort on side-channel attacks has been put into this device for the purpose of implementing effective countermeasures.

Randomization weakness

Even though the byte size required to generate cryptocurrency keys makes it statistically impossible to generate two identical keys, computers are notoriously unreliable at generating true entropy (randomness). A sufficiently dedicated attacker might be able to use weakness in software randomness to successfully replicate a private key,

Open-source solutions such as Glacier protocol solve this by introducing a source of true physical entropy, such as certified poker dices. This again is sufficient for an individual holder with infrequent demand for new keys, but unscalable for an enterprise in need of fast onboarding of new customers and receiving new deposits, especially if completely segregated wallets are required (i.e. not only segregated accounts in one HD wallet).

Other enterprise software solutions based on either multi-signature or multi-party computation can overcome this issue by integrating a more scalable source of entropy for the key generation but have to keep the integration channel between the source and the software generating the key secure to prevent man-in-the-middle attack.

This again is an area where even legacy HSMs stand out by their very design and original purpose because - by the industry best practice and certification requirements - they typically include multiple sources of hardware entropy, which are securely integrated with the key generation part within the module, preventing even a potential attack on the integration architecture.

Attack Vectors
Quantum Computing

Another potential weakness in current cryptographic algorithms comes from the promise of quantum computing, with which it might be possible to calculate the private key from the public key.

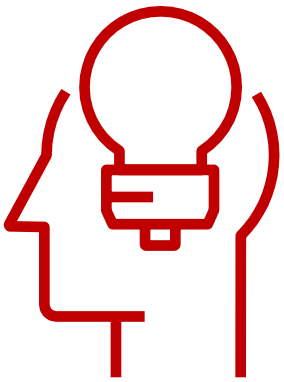

Cryptocurrency wallet software deals with this naturally by following the practice established with Bitcoin and the intention of its inventor(s) in using only an address, which is a (double-)hash of the public key related to the asset's private key. Since such software is typically used in all multi-signature and SMPC applications, they are naturally robust against this attack vector even though still being vulnerable to others.

Legacy HSMs, while protecting against physical, side-channel, and randomization vulnerabilities, however expose their operators to this type of attack because they require the operator to retrieve a potentially quantum-vulnerable public key in order for the business application to generate and use a respective address associated with the asset key.

Securosys HSMs deal with this potential vulnerability by allowing the developers to export only the address until after the first signature with the private key. This way, the keys created and retained in the HSMs don't have their public key exposed until the asset would have typically been withdrawn and thus the key has been rendered worthless.

Attack Vectors
Unauthorized Operations

Preventing unauthorized access to the key material is arguably less than half of the job when it comes to protecting crypto assets. The real challenge is controlling the use of the keys. Unlike in legacy applications, the damage from one-off compromising of a private cryptocurrency key can go to millions, if not hundreds of millions of dollars and is practically irreversible. The attacker might not even have a physical access to the keys – they merely need to exploit an application, which has been given permission to use the keys locally or remotely.

Authorization of transactions is at the heart of multi-signature - it's why the concept has been introduced in the first place. It is indisputable that it significantly increases security of the asset under such protection by introducing necessity to attack multiple keystores or authorized applications. It also - if correctly implemented - diminishes the risk of fraud to practically zero because it would require multiple fraudulent actors to collude.

There are however significant issues with multi-signature

- Not all cryptocurrencies support it.
- The implementations vary across different cryptocurrencies, adding to the complexity of transaction processing architecture.
- Because the transactions have larger bytesize, they also require higher fees to be paid.
- It doesn't allow for more complex groups and quorums.
- It is not possible to change the rules without moving the asset to a different address (some might argue this is a feature, but we see it as a shortcoming as long as it is not up to the custodian to define if they want these rules to be possible to change or not).
- The quorum requirements are exposed to the blockchain and thus subject to the same level of blockchain analysis as the asset itself.
- It doesn't allow for more sophisticated rules based on time, frequency, or volumes

Multi-party computation addresses most of these issues. Since it is performed on a single key required for the approved transaction, it is:

- Blockchain agnostic;
- As inexpensive as any other crypto transaction on a particular ledger;
- Allows for larger groups, quorums, and combinations thereof;
- Flexible on rule change;
- Private.

However, where it comes short is in two main areas:

- Its software-based implementation makes it vulnerable to the traditional attacks like malware, unsigned code execution, remote system access, keylogging, etc. While these attacks are made more complicated in linear proportion to the size of the quorum, they are still far from infeasible.
- It doesn't introduce the concept of time and other more sophisticated rules into the transaction authorization.

Legacy HSMs are pretty "dumb" machines in that they sign whatever they are requested to sign with an appropriate API authentication key. While this is sufficient for legacy PKI applications where a compromised key can be simply revoked and re-issued, it is unacceptable behavior for cryptocurrency applications.

This is the main reason why HSMs have had a pretty lukewarm reception by the cryptocurrency industry itself – CISOs know that the protection the HSMs focusing on physical and cryptographic vulnerabilities typically provide is insufficient.

This is the main reason why HSMs have had a pretty lukewarm reception by the cryptocurrency industry itself – CISOs know that the protection the HSMs focusing on physical and cryptographic vulnerabilities typically provide is insufficient.

It's also the reason why Securosys introduced several controlling mechanisms to help developers and security administrators to introduce additional controls over the private key operations. These controls include:

- Multi-party authorization allowing for flexible use of groups, quorums and combinations of thereof
- Time-locks
- Request timeouts
- Dedicated rules for key blocking, unblocking and even change of the rules themselves
- Verification of consistency of the rules through key attestation

In combination with the right design and operating procedures, such rules practically rule out any outsider or insider attack:

- Multi-party authorization and timeout might require n out of m financial officers to confirm a transaction within a certain time window. Combination of quorums might allow e.g. for a smaller consensus of board members to do the same.
- Time-locks could enforce a delay on any transaction before it is signed in order for an anti-fraud system or a monitoring team to kick in and raise an alarm
- Dedicated key blocking rules allow compliance and risk officers to block any transactions immediately if they suspect fraudulent behavior and unblocking rules might require large quorum of risk, security, and financial officers to allow the transactions to continue.

These rules are enforced within the same secure container, which physically protects the key material itself. It is impossible to sneak in an unsigned code or break the protection by a man-in-the-middle attack.

Attack Vectors
Failure Protection

Losing access to one's keys is as catastrophic as getting them stolen or compromised – it means that for all intents and purposes one loses the cryptocurrency balance.

Traditional multi-signature-based solutions introduce various methods of protection against the loss of key material. Glacier Protocol and similar concepts recommend using sufficiently low quorum-to-group ratio (e.g. 2/5) with geo-redundant paper wallets stored in physically protected locations. While this is arguably sufficient in protecting against physical damage, it is again very difficult to scale and useful only for individual holders. Other approaches might include more automated digital backup of the keys. However, all of these have one thing in common - they increase the risk of keys exposure linearly to the number of locations that either copies of the same key or multiple multi-signature keys are stored.

Because multi-party computation splits a single key, it is sufficient to simply back up the key material in case of a catastrophic loss of a required quorum. This however again introduced additional point of exposure of the key material, which has to be physically protected.

While top-of-the-shelf HSMs use data storage with much higher reliability standards than traditional servers or PCs, their potential failure still can't be ruled out and must be protected against by redundancy or regular backup. However, HSMs are also traditionally known for a very complicated and unscalable redundancy setup. Securosys HSMs address this by introducing seamless redundancy setup process, which includes highly secure physical initial pairing using master-slave model with up to 64 devices, real-time end-to-end encrypted synchronization

Attack Vectors
Regulatory and Market Considerations

One major area of consideration for many custodians, exchanges, and custodian platforms is how they are required to manage their customer holdings according to interpretations of various legislations around the world or by expectations of their potential customers. While we are admittedly no experts in the field, we understand that the interpretations in what constitutes ownership and custodianship of the assets vary not only across jurisdictions, but also across legal experts in the same jurisdiction. Additional level of complexity is added for enterprises with varying business models and their respective functional and legal requirements.

To meet these varying demands, the custodial platform must be flexible in its ability to decouple ownership of the key and the control of the asset if necessary.

That means in some jurisdictions or for some customers it might be necessary to have the key material stored on a device or premises owned and controlled by the custodian without the ability to use the key; while in other cases it will be necessary for the custodian to have the full control of the asset (i.e. the ability to use the key) while retaining legal ownership with the customer.

Offline paper-based solutions have limited flexibility in offering these various modes. Plaintext unencrypted private key or its representation on a QR code gives the full control to the custodian. A passphrase-encrypted private key or an HD wallet with a passphrase added to its mnemonic seed allows to decouple custodianship of the medium from the control of the asset it links to, but it comes with overhead - the customer would have to become part of the key creation ceremony to make the process feasible. At the same time, the customer would be at risk of losing the access completely by their own mistake.
Multi-signature - where supported - offers additional flexibility depending on how the quorum requirements are designed but comes with challenges described above.

And finally, this method might be operationally unsustainable if the competitive pressure requires the custodian to allow for more frequent deposits and withdrawals.

Multi-Party Computation is not too different to multi-signature in the options it provides but has the advantage of universality of its application and thus much lower operational overhead.

Legacy HSMs provide very little flexibility in this arrangement. Whoever is the administrator has typically full control over the private key operations even though technically they cannot get their hands on the key material itself. The only limited option for decoupling is in a multi-tenancy setup where the partitions remain stored in an HSM while their access is granted to a third party, but even that access can typically be easily manipulated with.

Securosys HSMs with their Smart Key Attributes capability provide the highest level of flexibility in the arrangement. Unlike SMPC, where the key is physically split into multiple shares and thus no one can technically be considered its custodian (if that is the requirement), the key material in the HSM is by design held by a single party (the HSM operator). At the same time, its control can be fully ceded to the customer, a 3rd party, or any combination thereof if that is the requirement.

Trade-offs
Trust

While from our analysis we see the correct design, implementation, and operations of Securosys HSM as a superior solution to cryptoassets safeguarding, nothing comes without trade-offs. We see that there are three main ones.

There are open-source implementations of both multi-signature and SMPC approaches, which makes their software layer more trustless than proprietary solutions. Furthermore, the flexibility to run them on various types of hardware gives the operators a chance to choose the most reputable, most open, most audited, or otherwise the most trusted combination of hardware components. However, looking at some of the most glaring examples of vulnerabilities in open-source software, such as Heartbleed in OpenSSL, and in mass-produced hardware like Intel architecture's Spectre, Meltdown and Plundervolt, it is clear, that open source cannot ascertain security.

HSMs are on the other end of the spectrum with their proprietary design and closed-source software. And while one can argue that economical and business incentives of their producers (including Securosys) can be equally strong or stronger in protection against intentional or accidental vulnerabilities as open source is, it is up to the evaluator to decide on how much value they want to put on each side of the decision scale. Securosys allows its customers a full audit of the software code and the design and manufacturing. In addition, its location in Switzerland gives the highest jurisdictional guarantee of resistance towards any pressures to lower our security design and operational standards.

Price

Since the open-source multi-signature and SMPC implementations start at no cost, it of course makes top-of-the-shelf HSMs with their production costs in thousands of dollars obviously much more expensive, for many startups prohibitively so.

For those who want to take advantage of the HSMs' security features while managing their expenditures, Securosys offers affordable HSM-as-a-service, which - with our Remote Partition Administration feature - doesn't require any trust in our management of the service in order to keep the access to the key material strictly controlled and private.

Trade-offs
User (Developer) Experience

Legacy HSMs are not famous for their great experience for neither the operators nor the developers. Purposefully built, software-only open- or closed-source implementations of the above-mentioned cryptocurrency safekeeping technologies on the other hand compete in delighting their administrators and developers and clearly come ahead in terms of simplicity and ease of use.

Securosys, understanding that our real competition are not legacy HSMs with their lackluster UX, but these cutting-edge software solutions strives to continuously improve the developer experience by natively supporting many blockchain- and cryptocurrency-specific requirements such as BIP32 or address export, by simplifying both administrator and application programming interfaces and by providing clear documentation and tutorials. To experience more, feel welcome to join our development program.

Final Notes

One important thing to note is, that these approaches are mutually not exclusive. One can benefit from hardware-based entropy and tamper protection of an HSM when using it in combination with multi-signature or SMPC implementation, as some of our customers who appreciate the blockchain-specific features of Securosys HSMs do.

Conclusion

We started this analysis trying to find out if our line of product is still relevant for cryptocurrency custodian applications, knowing that our main competitors are not legacy HSM manufacturers who didn't really pay attention to this market, but rather innovative trustless solutions like multi-signature and SMPC. We conclude it with confidence in our blockchain-focused HSM and in its ability can help secure cryptoassets compared to other solutions on the market. We also remain aware that this space is evolving so rapidly, that we have to continue to innovate and improve developer experience further to stay competitive. Finally, we hope, that this analysis - even though created for our own purposes - will help others make the best architectural decisions.

Any comments, feedback, or corrections of this material are more than welcome. If you have any, please reach out to productmanagment@securosys.com

To learn more about our products offering for both early stage startups to large crypto enterprises, please reach out to info@securosys.com or visit us on www.securosys.com

You can find our Developer's Program here.

securosys

Whitepaper Safeguarding of Crypto Assets

www.securosys.com