Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Data Privacy for Digital Asset Systems

Jillian Mascelli

2023-059

Please cite this paper as:
Mascelli, Jillian (2023). "Data Privacy for Digital Asset Systems," Finance and Economics Discussion Series 2023-059. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2023.059.

Data Privacy for Digital Asset Systems

Jillian Mascelli*

September 2023

Abstract

Data privacy in digital asset systems is of sustained importance to end users. However, there can be disconnect between an end user's expectations of privacy while using a digital asset payment system and the system's actual treatment of collected, stored, and used data. This paper provides foundational primer on data privacy alongside qualitative and technical assessments of various approaches to data privacy frameworks and strategies relevant to the early stages of a digital asset system's design. Analysis relies initially on an outlay of foundational data privacy concepts, including anonymity, confidentiality, and full disclosure, alongside three differing approaches to data privacy frameworks. Analysis finds that some concepts, such as a desire for "cash-like anonymity," are based on false underlying assumptions. The paper moves away from a likely unattainable standard of anonymity and instead focuses on a hybrid approach to data privacy, inclusive of Cavoukian's privacy-by-design and popular applications of privacy-by-policy. This hybrid approach is visualized with a technical comparison of privacy-enhancing technologies (PET) across architectural layers, detailing both popular and emerging PETs relevant to digital asset systems which prioritize a hybrid approach to confidentiality. The paper further finds that a particular combination of popular and emerging technologies may provide as-yet untested but novel benefits to maintaining strong confidentiality – and possibly end users' expectations of privacy – while data is under audit. A nuanced approach, rather than a reliance on a singular novel PET or dubious assurances of anonymity, may best facilitate strong confidentiality with sustainable end-user privacy protections for digital asset system users.

* The views expressed in this paper are solely those of the author and should not be interpreted as reflecting the views of the Board of Governors or the Federal Reserve System. The author would like to thank David Mills, Sonja Danburg, Nathan Palmer, Sarah Wright, Zach Vida, Cy Watsky, David Husband, Peter Lone, and Deidre Ryan of the Federal Reserve Board; the Federal Reserve Board's Technology Lab (TechLab) team; Katya Delak of the National Institute of Standards and Technology; Christopher Desch of the Federal Reserve Bank of New York; Erika Sales of the Federal Reserve Bank of Boston; and Angela Lawson, Alexander Lee, and Paul Wong for their contributions, reviews, and assistance towards this note. The author has previously published under the name Jillian Buttecali.

1.0 Introduction

This research note presents a high-level discussion of data privacy and the tools relevant to securing an end user's privacy while acquiring, holding, and using a digital asset, such as a cryptoasset. The paper presumes that robust end-user data privacy is a worthy precondition for the deployment of a digital asset payment system (Cheng, Wong, and Lawson, 2021, pg. 5). As such, analysis and conclusions below are most relevant to the design phase of a digital asset system prior to deployment.

Most digital payment systems employ some form of end-user privacy as a core feature. Not every digital payment system provides the same privacy protections or privacy features. Data privacy for a digital asset payment system depends upon which privacy-enhancing technologies (PET) and techniques are offered by the digital asset platform, and which of these tools an end user might be able, and willing, to employ. A digital asset privacy strategy could consider use case-specific protection of data collected, used, and stored by the system and incorporate a collection of privacy tools and techniques assembled to meet use case specific end-user privacy needs. This paper reviews possible approaches to implementing data privacy in a digital asset system, assesses several relevant digital asset PETs, and discusses architecture considerations when leveraging data privacy in a digital asset payment network.

2.0 Defining Data Privacy

For the purposes of this paper, data privacy is the assurance to a system user that the confidentiality, integrity, and availability to their information is protected (Barker, Smid, Branstad, and Chokhani, 2013). Also, for the purposes of this paper, digital assets are defined as digital objects that function as a thing of value, or a representation of a thing of value, to the asset holder in their digital wallet and are typically transferrable between holders.[1] Here, subsets of digital assets could include tokenized securities, tokenized deposits, and cryptocurrencies. Digital assets can also be called cryptoassets when underpinned by a cryptographic technology, such as a blockchain-based transaction ledger.

In a digital asset payment system, data privacy helps protect payment system users from unintended consequences of using a payment network, such as identity theft, discrimination, or public scrutiny of personal spending habits. It also helps protect payment system operators, who may have privileged access to user data, from operational, reputational, legal, and counterparty risks (Khan, 2018). End user concerns over data privacy, especially in the context of digital assets, are enduring.[2] Balancing risk mitigation

[1] The term is not used under the broader information systems definition for "digital asset," which can refer to a large swath of software, data, media, or documents.

[2] Analyzing public sentiment on electronic funds transfers (EFT) could help us better understand long-held user privacy concerns prior to the deployment of today's popular digital asset ecosystems, such as Ethereum, which are relatively new in the long

measures and end user privacy expectations can be a challenging task. Thus, a strategic approach to privacy within a digital asset payment system's design and surrounding architecture can help deliver on user expectations of data protection while mitigating risks faced by system operators and participants.

2.1 Approaches to Data Privacy Strategies

A digital asset system's design would likely consider how it protects privacy for data at rest, in transit, and in use by adhering to a data privacy strategy. A data privacy strategy, sometimes called a data privacy scheme, is the collection of tools and techniques used to create end-to-end data privacy, delivering on user expectations of data protection. A data privacy strategy consists of both tools and techniques used to create end-to-end data privacy while delivering on user expectations of data protection. There are two basic frameworks to support the development of a data privacy strategy: privacy-by-design and privacy-by-policy. While not mutually exclusive, understanding these distinct philosophies can be helpful to craft and implement an effective privacy strategy for a digital asset payment system.

Privacy-by-design is the proactive assessment of privacy needs given architecture choices, user requirements, and system data lifecycle expectations (Cavoukian, 2011). This approach would ensure that technical architects and product owners shape the system's architecture with privacy as a priority, performed in the earliest stages of the system design process. Privacy-by-design encourages end-to-end information management practices throughout a system's lifecycle, not just at a specific end user touchpoint, such as a login page or during cookie allocation. This approach to data privacy would apply to both a digital asset's physical and virtual infrastructure, IT systems, interfaces, and business practices. Additionally, within this methodology, only the minimum required user data would be collected, transmitted, and stored by the digital asset payment system. All other data would either remain uncollected or be collected and stored elsewhere as needed by the end user.

Privacy-by-policy is the implementation of business practices that promote the user's informed consent for an organization to hold, process, or transmit the user's data (Spiekermann, 2009). An organization could operationalize a policy-focused data privacy strategy for a digital asset system by offering transparent choices to the end user. A policy-based approach would likely not be developed by a system architect but, instead, would be within the operator's legal or business functions. Legal or business actors would likely assess the local jurisdiction's legal requirements and the primary business needs of the

history of payments. A 1982 study commissioned by the United States Senate concluded that EFT users were concerned about third-party insight into transactions and personal data. Users further wanted direct and detailed control over their own data. See also https://www.princeton.edu/~ota/disk3/1982/8223/8223.PDF page 11. Similarly, a recent survey conducted by the European Central Bank (ECB) concluded that users in their jurisdiction want novel digital asset payments "to remain a private matter" even if such privacy would restrict, say, a digital euro's features, such as third-party innovations and offline availability. See also https://www.ecb.europa.eu/pub/pdf/other/Eurosystem_report_on_the_public_consultation_on_a_digital_euro~539fa8cd8d.en.pdf#page=15 page 15.

system to develop a suite of documentation outlining the system's data usage practices. This might include a plain language privacy policy that requires the user's informed consent prior to data collection. This approach could also include an internal corporate policy limiting access to user data based on roles and responsibilities within the company. Privacy-by-policy for a digital asset could also include allowing only certain activities with the user interfaced based on the specific data a user has chosen to provide, or not provide. Unlike privacy-by-design, the policy approach is driven by the goals of business decisionmakers to protect users through information and consent rather than through technical means, such as a particular encryption methodology.[3]

Hybrid approaches, which combine elements of privacy-by-policy and privacy-by-design, are also an option when designing a digital asset payment system. Apart from a single-sided approach, some data privacy frameworks mesh design and policy approaches to implement comprehensive data protections. A "hybrid" approach could meet the complex needs of a digital asset system and help to mitigate operational risks that might be left open if an organization were to pursue only one privacy philosophy. One example of a hybrid framework is the Generally Accepted Privacy Principles (GAPP). GAPP is not specific to digital assets but is applicable a variety of financial system use cases. Originally developed by a consortium of accountants, GAPP dictates a series of data policy and documentation steps to properly protect the confidentiality of personally identifiable information (PII) (Johnson, 2009). This framework rests heavily on privacy-by-policy but also weaves in elements of privacy-by-design by outlining helpful PETs and proper data life-cycle activities to maintain privacy. Another hybrid framework option is the National Institute of Standards and Technology (NIST) Privacy Framework, which leans towards privacy-by-policy but includes a robust lexicon with which to develop a comprehensive and technical data privacy design.

2.2 Confidentiality and Anonymity

Data privacy encompasses a spectrum of related but distinct privacy concepts of anonymity, confidentiality, and full disclosure.[4] These three concepts are often confused to be synonymous. This paper holds that the opposite is true, that anonymity, confidentiality, and full disclosure are three distinct concepts within the larger spectrum of data privacy (Figure 1). Neither anonymity nor full disclosure are the primary topics of this paper. This paper focuses instead on various approaches specifically to confidentiality, such as privacy-

[3] While this is a valuable vantage point, this paper focuses on system design-based considerations for a technical privacy-by-design or "hybrid" approach.

[4] This research note lays out definitions for several important concepts to better understand the technical aspects of privacy. Some of these technical terms may be the same as legal terms. However, for the purposes of this paper, definitions of privacy, confidentiality, and other concepts are within a technical mindset, which may differ from a similar term's legal or policy definition.

enhancing technologies and techniques, to build a thoughtful system-wide privacy strategy. Let's explore the key aspects and differences between these three types of data privacy.

Figure 1: Spectrum of Data Privacy

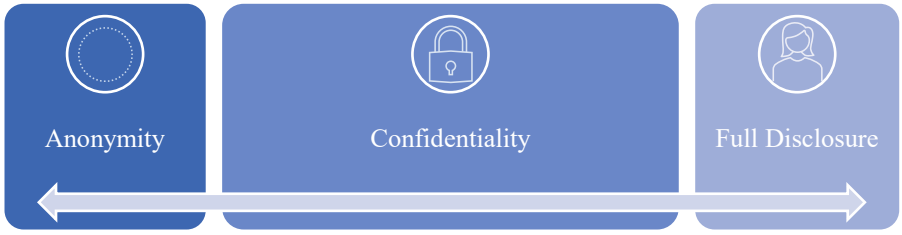

First, **anonymity** implies that data collected or stored by a system cannot uniquely identify an individual actor (National Institute of Standards and Technology, 2011). A digital asset system could facilitate anonymity by collecting less information about an end user so as not to be able to uniquely identify the user. Anonymity is nearly impossible to maintain in an electronic payment system that collects or stores identification data on end users. And anonymity is lost when a network actor can be recognized as a known identity instead of simply a distinct entity (Garfinkel, 2015). Anonymity may be difficult to maintain even in a digital asset payment system that does not require the collection of identification data. A payment system's underlying technology infrastructure will likely collect and log some data by default on every interaction with its system, such as logging a user's Internet Protocol (IP) address or details about the user's hardware when the user attempts to connect to the payment system. This default data collection may be mitigated using a distributed payment network in which each connected computer, or node, is responsible for keeping its own version of the network's operating software. However, as is common to public, distributed digital asset payment works, a shared transaction ledger recording pseudonymous identifiers (explained in the Tools section below) alongside transaction meta data. A shared ledger in a distributed digital asset network would likely prevent anonymity as thought of by NIST and Garfinkel – the user would be 1) uniquely identifiable and 2) known through heuristic analysis of transaction histories.

Second, **confidentiality** implies that collected and stored data is protected from view in some manner, such as obfuscation or access restriction, and available only to authorized actors.[5] Varying degrees of data confidentiality fall between the extremes of anonymity and full disclosure. An architect could design a digital asset system that masks or prevents access to end user data collected by the system. In that case, this paper considers such obfuscation as *confidentiality*, not anonymity. Further, all following sections within this paper discuss technologies and techniques relevant to confidentiality, rather than those that would ensure anonymity.

[5] See ISO documentation for ISO/IEC 27000:2018, section 3.10. https://www.iso.org/obp/ui/#iso:std:iso-iec:27000:ed-5:v1:en

Third, **full disclosure** implies that collected and stored data is not protected from view by any system user. Full disclosure exists on the opposite extreme of anonymity in the data privacy spectrum (Figure 1). Confidentiality and full disclosure are two distinct concepts. As such, fully disclosed data are no longer confidential. Privacy remains relevant, however, even to fully disclosed data as a system might allow a user to later obfuscate or remove fully disclosed data. Depending on the flexibility of a digital asset system's design, one use case for full disclosure of user data could be when a user selects to share their payment transaction data, including payer and payee digital wallet addresses (as described in the privacy tools section below), transaction value, or associated metadata, with a public distributed ledger network. Here, it is possible that some data points within the same transaction could be considered confidential, such as collected location data for the payer, while other data could be considered fully disclosed to all system users, such as the payment's value.

2.3 "Cash-like" Privacy

A robust data privacy strategy should not confuse anonymity in digital asset payments with "cash-like" anonymity acquired through physical banknote transactions. Certainly, there are notable similarities between physical banknotes and some digital assets, such as facilitating peer-to-peer payments without a third-party intermediary. Cash transactions could also be seen as akin to token-based transactions common in some cryptoasset marketplaces rather than account-based transactions common in online retail banking. However, using the appellation "asset" for digital assets may be misleading here. Further, the concept of "true anonymity" for an asset, even one designed to function like a virtual banknote withdrawable from a virtual ATM, often takes an overly simplified approach to the realities of how modern networks collect connection, logging, hardware, and other data points by default.[6]

Comparing the privacy inherent to physical banknotes or bearer instruments with data privacy for digital assets is a false equivalency. Cash is an item exchanged in physical form; cash requires no ledger of transactions to complete settlement; and cash presents few barriers to use other than physical possession. A cash payment requires only one data point to clear and settle: value of the transaction. Further, a cash payment does not require any hardware or software connection to an information system, database, or computer network. Conversely, to clear and settle a transaction, digital assets collect more data than simply the value (Figure 2). Many considerations, such as unauthorized access to payment data and identity theft, are relevant to the privacy strategy of a digital asset payment system but not relevant to a banknote

[6] Here, I am questioning the theory presented by Goodell et al., pg. 18, in reference to "true anonymity" and cash-like properties an end user can (or should) expect when a digital asset system collects a minimum amount of the user's data to support a transaction. A tenured system administrator or security researcher could likely uncover data unintentionally collected by or provided to such a system that would not, relying on NIST definitions outlined in this paper, qualify as anonymous. I also question the definition of "true" anonymity, as this term is not well defined in supporting literature elsewhere or in data privacy standards documentation.

transaction. Figure 2 illustrates this point by outlining the basic data required, collected, used, or stored from transactions occurring with cash, permissioned digital asset networks, and even permissionless digital asset networks, which some may consider as requiring less user data than a permissioned system.[7]

Figure 2: Transaction Data Requirement Comparisons

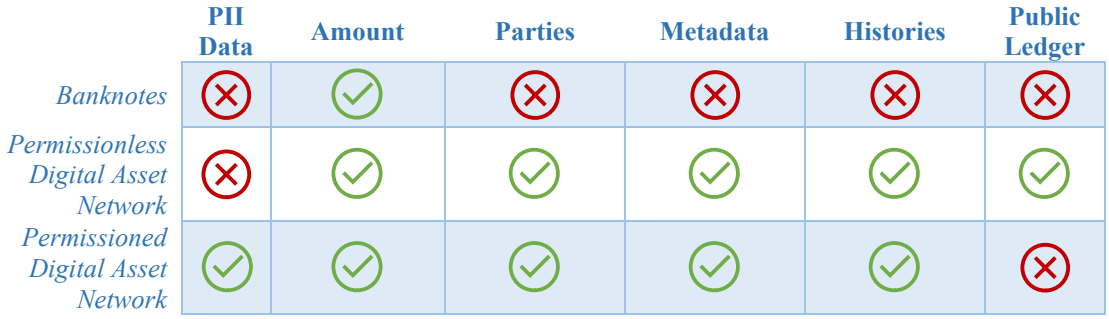

For these reasons, cash transactions should not be held as the standard by which confidentiality, or even anonymity, can be achieved in a digital asset system. Data privacy in a digital asset payment system should be instead tailored to the system's design requirements and surrounding architecture. This tailoring should be conducted through the development of a data privacy strategy rather than to reaching for "cash-like" anonymity.

3.0 Data Privacy Tools for Digital Assets

Many privacy technologies and techniques have developed both independently of and in tandem with digital payment innovations and the expansion of privacy regulations over time.[8] However, this section looks at a limited range of current privacy technologies and techniques. The curated selection below includes both popular and emerging technology options relevant to data privacy in digital asset payment systems.[9] More

[7] In Figure 2, an "X" indicates not applicable, and a check mark indicates applicability. While all marks are subjective, I rely here on my previous six years in the Federal Reserve Board's banknote section, my current work in the same organization's TechLab payments technology research team, the research and analysis that went into developing this note and its collection of references, and feedback from colleagues with additional expertise in digital asset technologies.

[8] Advancements in data privacy technologies became increasingly relevant to digital payments after the passage of the Electronic Funds Transfer Act of 1974 (EFTA), which codified consumer protections for emerging payment options. Instead, the architecture presented here and elsewhere in other digital currency wallet designs proposed as "cash-like" might better be defined as systems supporting multiple, redundant layers of confidentiality rather than anonymity as defined by NIST and Section 2.

[9] Implementing the proper balance of privacy-enhancing features is challenging. The PETs described in this paper each facilitate anonymity or confidentiality, but to different ends and for different purposes. Some data privacy topics, such as synthetic data through machine learning, were deemed out-of-scope for this paper, which is intended as a broad look across data privacy topics generally relevant for digital currencies. Not every data privacy technology is relevant for every digital asset design, however, and those included in this paper are only discussed in basic detail. The technologies chosen for a digital asset system should be appropriate for the specific data collected and stored by the system while still allowing the issuer to meet end user requirements. Pseudonyms, for example, provide privacy only for a user's identity, but asymmetric key pair cryptography provides a layer of confidentiality for a user's transaction history, data in transit, and data at rest. Table 1 details which of the privacy tools explored in this paper are most relevant to specific user actions.

extensive lists can certainly be found elsewhere in the expanding academic literature on data privacy for digital asset, and other, payment systems.

This section first covers several privacy tools, some common and some novel (in this paper "novel" does not necessarily mean new but indicates an innovative option not yet deployed at scale in digital payment systems). These tools can assist in the implementation of a data privacy strategy after a system designer or architect identifies the privacy framework that best meets their needs. This selection is tailored towards those technologies that are particularly relevant to digital asset networks running a native asset, with multiple participants, and whose primary purpose is to facilitate asset transactions. The section then assesses how and when certain technologies might protect the confidentiality of data collected, stored, and/or in use by a digital asset payment system.

3.1 Foundational Privacy-Enhancing Technologies (PET)

Multiple PETs can be included in the design of a digital asset payment system. These technologies are often platform agnostic (for example, applicable to blockchain ledgers or relational database systems) and invisible to end users interacting with an interface.[10] PETs can have wide or limited applicability for digital payments depending on how a system uses these technologies. An assessment of the data collected, used, and stored in the digital asset system, and a deep understanding of the privacy use cases for that data, would help a system designer determine the relevant privacy tools to protect data confidentiality throughout the system.

Encryption is a fundamental technology used to protect data confidentiality, both in transit and at rest. Data encryption is the method by which information is converted from its original form into ciphertext, which can be viewed as blocks of seemingly random alphanumeric characters. Encryption, a commonly used technology for privacy protection, is extensible and customizable but typically meant for temporary, rather than permanent, protection from view (Asrow and Samonas, 2021, pg. 6). It can limit visibility of user data in all use cases or only specific uses. If encryption is in use to protect data in transit and/or at rest, the digital payment system would use a cryptographic algorithm and cryptographic "keys" to convert the original information into something called ciphertext, prior to data transmission. The original information would be processed through a series of computations and the resulting encrypted data could then be used, transmitted, or stored while remaining confidential. A data sender and data recipient can keep the encrypted data protected while sending it back and forth, revealing the original message only when needed. The data's possessor can decode the information back into its original readable form or compare the encrypted data

[10] This paper does not assess every potential privacy-enhancing technology, tool, or platform specifically relevant to a digital asset. Instead, the options presented here are a sampling of relevant PETs to a starting point for data privacy professionals developing a digital asset privacy strategy.

with the ciphertext they expected to receive. By revealing the original message or comparing the expected information with the received information, a payment recipient can confirm the data's validity. Encryption supports a variety of digital asset payment use cases, such as protecting transaction data from unauthorized access.

Asymmetric key pair cryptography keeps data confidential through encryption while the data is at rest or in transit.[11] (Here, "at rest" means data being stored within the system but not moving between multiple system users.) This approach uses one private file of alphanumeric characters (the private key) and one public file of alphanumeric characters (the public key), and each key file can either block or unblock data from view. Typically, the private key is known only to the original owner while the public key can be provided to multiple parties.

Separately, **symmetric key cryptography** uses only one private key to encrypt and decrypt data at rest. This method is faster but less secure than asymmetric key pair cryptography because the private key is shared between two or more parties. Notably, possessing the private key does not prove that an actor is the rightful owner. If a bad actor presented the correct key, they could convincingly pose as someone else in a transaction and siphon digital asset funds away from the true owner. One way to protect a key's confidentiality is with digital signatures.

Digital signatures use asymmetric key pair cryptography (a private key and public key pair) to allow data verification without fully disclosing the private key. A digital signature is a hash of some content, expressed as a set of numbers. This signature proves a private key's validity without unmasking the key. Similarly, a digital signature helps to prove that an individual is the owner of some amount of digital asset. An actor can verify another actor's digital signature if they have the other's public key. Digital signatures are widely applicable to a variety of digital payment systems, so this tool is available for use with or without a blockchain or distributed ledger system.

Ring signatures employs digital signatures to allow users to obscure their identities by forming sets of anonymous, or theoretically anonymous, actors.[12] A set, or ring, of collected signatures helps mask which users are engaged in a specific transaction. While this is a compelling privacy feature, studies have shown that systems relying on ring signatures may still be vulnerable to privacy degradation through heuristic transaction data analysis (Möser et al, 2018, pgs. 4, 8 – 9).

A **payment address** is a digital location identifier for some store of digital asset value, expressed as a unique string of alphanumeric characters. Payment addresses rely on encryption to keep some data

[11] Barker and Dang (2015) provide a helpful review of key pair cryptography for both asymmetric and symmetric key cryptography and digital signatures. I particularly recommend their Glossary section (pgs. 78 – 84) for a plain language approach to terminology. I relied on it heavily throughout this section in addition to my previous work using public and private keys during various application development projects.

[12] Ring signatures are a key feature of the privacy coin Monero and are described here how they are deployed on that production platform. See also: https://www.getmonero.org/resources/moneropedia/ringsignatures.html

points confidential. A transaction counterparty can specify a payment address as a value transfer's endpoint. Digital asset systems might offer users different types of addresses from which a user can choose, such as single use, shielded, or transparent addresses.

When used only once, a single-use address, or a one-time destination address, makes it difficult for another actor to trace the receiver's transaction history outside of that single transaction. However, a user could send value to that address or receive value at that address multiple times. A user may also need to combine funds from multiple one-time addresses to have sufficient funds to complete a transaction, thereby allowing tracing services to link those addresses to the same owner. Using the address multiple times may be more convenient but erodes the feature's privacy function. Separately, a **shielded address** masks a transaction's sender, receiver, memo, and amount. Generally, actors outside of a transaction cannot view a shielded address.[13] A **transparent address** also obfuscates a transaction's participants, but the address is viewable outside of the transaction. One transaction participant could opt for a different type of address than their counterparty. The address choices partially determine the confidentiality of transaction's data after settlement. If all transaction participants use a transparent address, for example, then the transaction's value can be viewed by actors outside of the transaction.

Pseudonyms are unique identifiers of random alphanumeric characters, often generated by a computer system and provided to the user. A pseudonym does not mask, say, sensitive transaction data. Instead, a pseudonym, functioning as a payment address, is a unique identifier for a transaction participant instead of another identifier, such as a government identification number or legal name. A pseudonym's uniqueness combined with a history of payment transactions to and from that address prevent the identifier from maintaining true anonymity within the payment system. Instead, the pseudonymous identifier can help maintain the user's confidentiality, especially if the user employs more than one pseudonym for their activities within the system. Theoretically, such activities could mask the user's legal identity. However, because a system could allow administrators or users to look up a pseudonym's previous transactions, pseudonyms are vulnerable to data mining if a user transacts with their pseudonym more than once (Reid and Harrigan, 2013, pgs. 197-223). This facilitates forensic and behavioral data analysis, revealing links between system users to unmask personal spending habits and, in extreme cases, a user's real-world identity.[14] Such connections might be unintentional and unknown to the pseudonym's user.

[13] Shielded addresses and, as reviewed next, transparent addresses are options available to users of the privacy coin platform Zcash (ZEC). Zcash's documentation provides a brief explanation on the differences and user options. See: https://z.cash/learn/what-is-the-difference-between-shielded-and-transparent-zcash

[14] The blockchain analytics company Chainalysis uses techniques like clustering to de-anonymize cryptoasset transactions on pseudonymous networks and cryptoasset exchanges. See: https://www.chainalysis.com

3.2 Emerging Options for Privacy Preservation during Data Auditing

Putting aside privacy models and focusing on novel PETs, secure multi-party computations, fully homomorphic encryption, and zero-knowledge proofs are three emerging technologies which could be employed as PETs in a digital asset system. These three options each offer a novel ability over, or in addition to, the PETs described in the previous section for privacy protection while data is at rest, in storage, and/or in use.

Secure multi-party computations (MPC) are cryptographic techniques that keep data confidential while an actor performs analysis on a data set. Theoretically, MPC could protect transaction data privacy during quantitative analysis on large data sets. While secure MPC is a less-explored tool for confidentiality protection in digital assets than other options in this paper, it is nonetheless an emerging option for privacy protection during auditing. To this end, MPC may be a critical tool in the balance between data privacy preservation for individuals and data auditing by a third party (Hastings et al, pg. 1, 2019).

Similarly, **fully homomorphic encryption** can keep transaction data confidential while in use by a third-party actor with a variety of benefits when paired with other tools, such as improving secure MPC's computational burden (Gentry, 2009, pgs. 4, 23 – 24, 35). A third-party actor uses a private key to examine the encrypted transaction data for analysis without first decrypting the protected information. Fully homomorphic encryption may assist auditing processes while maintaining confidentiality for transaction participants and system users. For digital retail transactions, fully homomorphic encryption could also add a layer of data protection for the consumer in a payment system where arbitrarily created keypairs serve as account identifiers.

Lastly, **Zero knowledge proofs (ZKP)** ingest transaction data from the sender and output an assessment on whether the provided data is collectively true or false. A ZKP is not a mathematical proof but a clever use of cryptography. ZKPs obscure the provided data from anyone but the originator. A ZKP is a cryptographic tool not controlled by a human actor, mitigating the risk of human error or malicious intent. The data recipient does not personally inspect the sender's data but will rely on the ZKP's automated determination that the sender is the correct individual and possesses enough value to complete a transaction. ZKPs are resistant to data mining or forensic analysis, even with a shared ledger. The ZKPs can help keep a sender's digital asset holdings confidential from the entire payment network. The two forms of ZKPs most relevant to digital assets are Succinct Non-interactive Arguments of Knowledge (zk-SNARKs), which are more computationally complex, and Scalable Transparent Arguments of Knowledge (zk-STARKs), which consume more memory.[15] zk-STARKs are deployed in only a few notable public digital asset

[15] See https://docs.ethhub.io/ethereum-roadmap/layer-2-scaling/zk-starks/ for Ethereum's discussion of zk-STARKS in their production ecosystem.

platforms. Computational intensity and operational costs are key detractors to selecting a ZKP as a system's privacy technology.

Malicious actors can use digital asset platforms to avoid detection and identification by law enforcement. For the purposes of facilitating privacy in a digital asset, however, anti-money laundering (AML) and similar efforts could be primarily addressed by using features that facilitate data confidentiality while providing some level of auditability, such as a combination of multi-party computations, fully homomorphic encryption, and ZKPs for protecting confidentiality during auditing or during use in data analysis. Further, a digital asset payment system's issuing layer could encode AML and auditing with tools like smart contracts that define transaction limits, automating privacy policies with code. Adhering to jurisdictional requirements may adjust the required composition of privacy tools at the issuing layer.

Additionally, data privacy analysis models paired with a novel PET may assist with minimizing the risks of disclosure while an actor performs large data set analysis. Privacy models, such as local differential privacy, limit the risks of analyzing data by keeping the data set within specific conditions. When layering in a novel technology with a privacy model during large data set analysis, additional privacy benefits may be available to data analysts or auditors without disclosing or revealing personal data, such as applying fully homographic encryption to conduct data analysis without decrypting the data while in use. Such approaches could allow for efficient and automated data analysis and auditing on a system without some of the processes currently in use to manually (and often laboriously) protect user privacy. Combining a privacy model and one or more novel PET during analysis may also, when automated, help prevent human error during analysis and auditing. However, more research would be needed here to prove out which combination(s) of modeling techniques and novel PETs. Currently, some tested combinations that would best protect privacy in large data fail to facilitate useful analysis, leading analysts to fall back to more traditional, and manual, methods (Davidow and Manevitch, 2023).

3.3 Data Privacy Techniques

Data privacy techniques are the capabilities of an information system to keep confidential a user's data while balancing others' needs to access that same data. Privacy techniques may employ one or several technologies and policies in their approach to protecting confidentiality. Access procedures, collection limitations, retention requirements, redaction, and federated databasing are all data privacy techniques broadly applicable to end user confidentiality in digital payment systems. Such techniques typically allow for both auditability and confidentiality, no matter the underlying technology in place. Along the "privacy spectrum," access to a user's information can be unlimited (full disclosure) or limited (confidentiality). For

limited data access, several techniques can guide a digital asset system designer to tailor their approach to data privacy protections. Three common approaches are:

- Discovery Access – Limited access is granted to allow the system to reveal types, summarized or select portions of the data without revealing the entire dataset
- Role-Based Access – Different roles can be granted different levels of access
- Time-Based Access/Data Leasing – Data Access can be scope for a limited time

Apart from access controls, records management is also a helpful technique to ensure confidentiality. Records management approaches include:

- Collection limitations – The system doesn't collect more than it needs to operate
- Retention limitations – The system doesn't store data for longer than required. Either purging records in part or in whole automatically via a defined process.

Lastly, another foundational privacy technique is to limit the exposed data shown to others, even when legitimately accessed by an approved system actor. Such limitations in visibility include:

- Selective revelation – An actor must request to view a specifically scoped set(s) of data and provide justification, which is then reviewed by another party. If approved, a system authority may reveal this selected data set to the requestor with or without the consent of whom the data is about but with consent of the information owner or steward. The justification for each revelation should be auditable.
- Selective redaction – A data request can be broad while automatically excluding certain data fields or selectors. An example of this form of redaction is the automatic detection of license plates or faces by an approved data collector, who then blurs the resulting images meant for public searches.
- Federated Databases – Data isn't centrally stored or managed but can be selectively discovered or revealed as applicable by law or policy to allow data feeds to be orchestrated. Combined with other privacy protecting tools, suspected illicit payment flows could be tracked in an intermediated system without breaching the trust and confidentiality of activity that isn't implicated. Standards for data models and ledgers, while beyond the scope of this paper, could also be implicated in a federated database design.

3.4 Digital Asset Privacy Tools

Beyond individual privacy-enabling tools and techniques, a digital asset payment system can also leverage products unique to digital assets that can bundle and deploy a collection of various privacy tools. These

products with bundled privacy-focused technologies include tailored sets of PETs and privacy techniques to provide a sphere of data confidentiality for digital asset users. Because of the strength of these products, both good actors and bad could benefit from their privacy-preserving features. This paper does not explore the relative legality or legitimacy of privacy tools, simply their technical construction.

Privacy coins are a type of digital asset that integrates data protection methodologies and technologies as the preeminent design principle. They tend to collect less data at the outset and, even though the ledger may still be public like the Bitcoin ledger, privacy coins layer multiple, even redundant, PETs to keep user data confidential before, during, and after a transaction to an extent that other public ledgers do not. Currently, the two main privacy coins available to the public are Monero and Zcash. The confidentiality provided by privacy coins could be described as cooperative. These systems are designed to protect the network's data by enabling privacy tools or techniques that equally apply and protect a group of users. Shielded addresses, ZKPs, encryption, and other tools are concurrently leveraged by a privacy coin user to redundantly protect the same data points. However, the degree to which privacy coins shield transaction data in practice fluctuates due to the optional nature of some privacy coin features. For example, a privacy coin network may offer transacting parties the option of shielded addresses, but if the network does not compel the use of this privacy tool than the user isn't guaranteed its protection unless they select the option. Some users may opt for better transaction speeds instead of leveraging all the PETs available to them with a privacy coin. Or a user may use an exchange to complete a privacy coin transaction. In those cases, the user may not be realizing the intended confidentiality enabled by the privacy coin but potentially limited by an exchange.

Smart contracts, like privacy coins, are also not a singular tool. This self-contained codebase dictates how a network should process a transaction and can attempt anonymity by limiting the user data collected to participate in a payment. A smart contract could support confidentiality by protecting data that the transaction network collects from the user. More ambitious smart contract protocols, such as the Aztec Network developed on Ethereum, aim to provide a foundational level of privacy to create a group of fully private transactions safeguarded by zero-knowledge proofs.[16] Smart contracts, however, are not invulnerable to risks such as fraud or human error, even when the contract's code is properly executed.

Mixing services, or mixers, bundle deposited information into a black box of encrypted storage, where they can be withdrawn by another individual with the correct cryptographic keys.[17] Mixers mask the link between a deposit and withdrawal by bundling this data with many other deposits and withdrawals, called an anonymity pool, where combined data from multiple users provides mutual cover, or collective

[16] For more background on the Aztec Network, see https://aztec.network/.
[17] Mixers can be created with a variety of cryptographic tools, including smart contracts leveraging zero knowledge proofs. However, mixing services have facilitated financial crimes including money laundering by leveraging mixers as unregistered money services businesses (MSB).

privacy, for each other's transactions. Rather than simply mask a particular data point, mixers attempt to anonymize relationships altogether.[18]

3.5 Privacy Tool Considerations for Tiered Architectures

The distribution of a digital asset from an issuer to an asset holder can be visualized within functional layers or tiers.[19] In production digital asset ecosystems, assets flow in a complex web of devices, networks, and actors. In this paper, the term "tiers" generalizes this complexity to refer to the distribution architecture for a digital asset from its issuer(s) to distributors and end users.

A single-tiered distribution facilitates a digital asset issuer to directly mint, issue, and distribute an asset to asset holders within the same network and without intermediaries. A multi-tiered distribution model, however, implies that there are friction points, network differences, and intermediaries involved in distributing the asset from the issuer to an end user holding and interacting with the asset. This paper focuses on the multi-tiered distribution model, instead of the single-tiered model, to broaden the data privacy topics available for discussion. This paper does not hold an opinion on which model is superior for distribution efficacy or data privacy protection. Here, conceptualizing tiers is simply a mechanism for a deeper conversation on data privacy considerations throughout the diverse possible use cases in a somewhat complex cryptoasset ecosystem.

A basic, multi-tiered digital asset could be distilled into two components: the issuance tier and the circulation tier (Figure 3). Under this two-tier distribution model, the issuance tier and circulation tier would have separate pools of users, data, and use cases, implying that issuance and circulation require different privacy considerations. Conceptualizing the distribution of digital asset as tiers helps determine which privacy technologies and techniques should be employed for specific issuance activities or circulation

[18] Mixers and tumblers could lead to laundering-type activities and can be subject to the Bank Secrecy Act. See https://www.fincen.gov/sites/default/files/shared/FIN-2013-G001.pdf for 2013 guidance from the Department of the Treasury Financial Crimes Enforcement Network.

[19] Bindseil (2020, pgs.2, 19, 20) conceptualizes a multi-tiered central bank digital currency (CBDC) arrangement as a possible policy hedge against a CBDC's potential risk facilitating both disintermediation and bank runs. His concept of "tiers" relies less on the information technology (IT) concept of physically distinct infrastructure and hardware tiers and, instead, relies on the concept of tiered remuneration where designated actors can control the circulation of CBDC funds in circulation. I present no opinions on this proposal. Instead, my use of the term "tiers" for digital assets is closer to the traditional IT use of the term, where funds flow from an issuance tier and into a distinctly different circulation tier, which may include a different set of actors, devices, network(s), and geographical locations. This is meant to encourage the reader to "cast a wide net" when thinking about a digital asset system's data privacy strategy to align with the realities of today's ever-expanding digital asset marketplace. Ideally, this section helps the reader to conceptualize the diverse privacy needs and implications of interconnected networks, devices, and actors which may all have different use cases and privacy needs for a digital currency. Common cryptocurrency parlance might instead separate these concepts into "layers," such as the blockchain at Layer 1 and distributed application services at Layer 2. For visualizations of Layer 1 and Layer 2 in use with a production digital asset ecosystem, see Ethereum's documentation: https://ethereum.org/en/layer-2/

activities. A specific PET, such as digital signatures, might not offer relevant privacy protections for every use case across both the issuance and circulation tiers (Table 1).

Figure 3: Digital Asset Tiers as Functional Layers

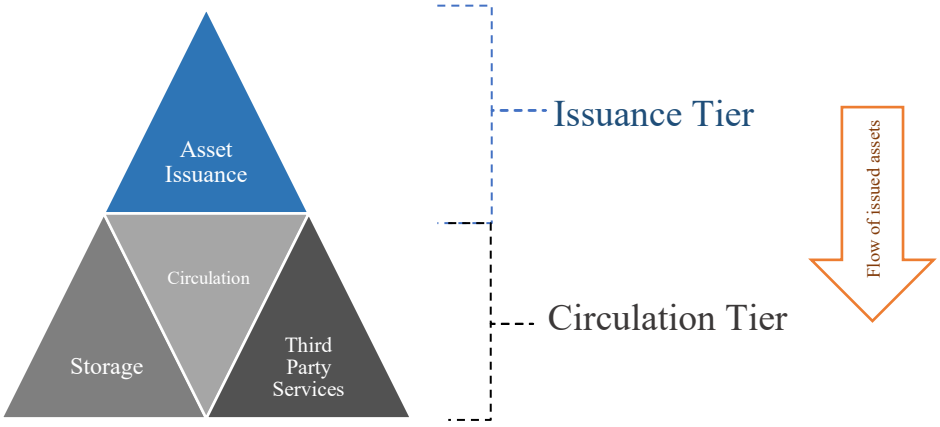

The issuance tier, comprised of nodes connected to the payment network that is responsible for generating and minting the digital asset, relies on intermediary parties in a separate circulation tier to distribute the asset to end users and provide enhanced functionalities. The issuance tier encompasses PETs relevant to the creation and initial distribution of the asset. The issuance tier focuses on the core minting and settlement platform tasked with infusing new assets (be them tokens, cryptoassets, or some other form) into the larger digital payment network. But, as these activities are limited and the parties potentially involved in these activities also limited, the data touchpoints involved in this base layer are likely easy to locate and account for in a data privacy strategy. By contrast, the circulation layer's data touchpoints are potentially much more expansive and could require a more nuanced approach.

The circulation tier's primary functionality is to take issued digital asset and distribute it to a wider set of users, such as businesses, retail banks, and consumers. The circulation layer for a multi-tier digital asset encompasses everything not specifically included in the acts of minting (or asset creation) and issuance into a general liquidity pool or network. Circulation activities might include online and offline transactions and transfers to an end user's wallet. The circulation layer encompasses PETs relevant to asset usage by an end user, rather than the issuer, such as storage in a digital wallet or use in a third-party service. This tier could include numerous data touchpoints, possibly more than the issuance tier. Touchpoints involved in circulation would be as onboarding applications, wallets, exchanges, and hardware. An end user application, for example, could include wallets and exchanges and provide a user interface for individuals and organizations to use to acquire, exchange, and store their digital asset value.

Between these two functional layers, many potential pieces of data are collected, formed, transmitted, and stored by digital asset payment systems and networks. Further, the circulation and issuance layers may or may not run within the same computer networks.[20] Depending on the network's design, key payment functions, such as transaction settlement, could even be spread across both the issuance and circulation layers. The circulation layer's systems and applications could leverage many of the same PETs and privacy techniques as the issuance layer, however in different ways. Because of their differences, each tier might require its own privacy strategy, one for issuance activities and one for circulation activities, leveraging a tailored collection of privacy tools to enhance data confidentiality across the digital currency's ecosystem.

3.6 Privacy Tool Use Cases for a Multi-Tiered System

Table 1 visualizes a simple assessment of which uses cases each PET and privacy technique explored in this paper might be best for data privacy protection. While not exhaustive, this visualization is a starting point for a digital asset system's designer to consider while developing a data privacy strategy. The designer might want to consider which data use cases are of highest priority for their digital asset's approach to privacy and which technologies and techniques might be well suited for the strategy. In developing a comprehensive data privacy strategy, a digital asset system designer might also consider which privacy tool would best accomplish the priority use cases' main user tasks. This connection between tools and tasks is visualized in table 1 with a shield and check mark.[21] The tools shown in Table 1 are each available to both layers. However, some may be judged as unnecessary for a particular use case. For example, a digital asset system leveraged by a small population of authenticated users could require shielded addresses on the issuing layer but not require shielded addresses for peer-to-peer payments within the circulation layer. The determination of which privacy tools are most relevant to a digital asset's tiers and use cases is subjective in Table 1. In this visualization, only those PETs, as laid out in earlier sections, judged to be most relevant to the corresponding use cases are marked. Additionally, for visual clarity, table 1 limits green check mark indicators to each tool's most applicable benefits (rather than an exhaustive list of each tool's possible benefits). Digital asset design choices could certainly alter this balance of helpfulness and relevance between tools and use cases.

[20] Depending on the degree of distributed control, access, and circulation, these layers could exist within wholly separate but interoperable computer networks. While important to note, decentralization is not an issue this paper seeks to address.
[21] Table 1 visualizes privacy tools using a general understanding of each tool's likely vulnerabilities and strengths and on how each tool is commonly used in existing digital asset networks, such as Bitcoin and Ethereum.

Table 1: Possible Digital Asset Uses and PET Matrix (Privacy-by-Design Approach)

	Prove Identity	Prove Sufficient Funds	Transmit Data	Store Data	Use Data	Audit Data
Symmetric Cryptography			✓	✓		
Asymmetric cryptography	✓		✓	✓		
Fully Homomorphic Encryption					✓	✓
Secure MPC						✓
Digital Signature	✓	✓				✓
Ring Signature				✓		
Pseudonym	✓					
Shielded Address			✓	✓		
Transparent Address			✓			✓
One-time Address			✓	✓		
Zero Knowledge Proof	✓	✓	✓	✓		✓
Transaction Mixer			✓	✓		

4.0 Additional Data Privacy Strategy Considerations

Technologies such as the ones described above could be implemented for a digital asset to facilitate data privacy-by-design. Several considerations would help a digital asset to meet a jurisdiction's privacy requirements and the end user's preferences. These considerations are the issuing layer's network architecture, selective disclosure, performance benchmarks (e.g., minimum throughput), and auditing needs.

4.1 Network Architectures

A digital asset network's architecture includes the software, hardware, and middleware resources needed to bring the network online and facilitate the transfer of funds. This architecture necessarily includes multiple data touchpoints, such as a user interface with a data collection form, servers which collect data on each connection, and transaction history ledgers that store payment data for some defined amount of time (i.e., a company's data retention policy requiring employees to store data collected from an online

form for a minimum of three years). Each data touchpoint is an additional area to consider privacy—whether to facilitate anonymity or protect confidentiality. Overall, data privacy choices are dependent on the network architecture's design. The network architecture determines the locations of data touchpoints and what types of data will be processed by the network. These architecture choices include whether the network is permissioned or permissionless and operates with a centralized database or distributed ledger. The architecture's design can combine the concepts described throughout the paper – confidentiality, PETs, and functional layers.

A **permissionless** digital asset payment network is publicly available for use without a central authority, such as Bitcoin. A **permissioned** payment network requires credentials from an authority to access and use the network and, possibly, to gain advanced abilities once connected to the network. The permissioned structure requires a central authority or an authority consortium that is trusted by transaction participants to administer network access and permissions. In this way, permissioned payment networks could resemble the existing financial rails for online banking. If so, for a permissioned digital asset network, the issuing layer's network access points may require some form of user identification to participate in the payment system. Data collection, then, would begin when an actor provided their identification data to obtain permission to use the digital asset network within the issuing layer. Privacy in permissioned networks may focus on employing confidentiality techniques of obfuscation and access restriction. We should not assume, however, that data privacy is more important for either a permissioned or permissionless network's participants. Instead, as discussed with the concept of privacy-by-design, the privacy strategy developed for a digital asset network should be relevant to and tailored to the network's architecture choices, the specific data collected by the system, and users' considerations.

Centralization of a digital asset network refers to the concentration of network control and typically implies that a central authority created, maintains, and likely controls access to the network. A network can have centralized properties even if the network supports distributed data storage. For example, distributed ledger technologies (DLT) like blockchains can operate within a centralized governance but have distributed data storage. A centralized payment network, even one operating a DLT, can support efficient implementations of novel privacy technologies.[22]

Decentralization implies that the network may have been created by one or many entities but is maintained by many entities distributed across the network. Decentralization and distribution should not be confused. (De)centralization refers to the balance of network control. Distribution refers to the dispersion

[22] For example, Corda, a permissioned DLT platform that can be configured to support either a centralized or decentralized architecture, has been shown through limited research to also support the application of zero knowledge proofs. See: https://www.coindesk.com/ing-bank-devises-privacy-fix-for-r3s-corda-blockchain

of actors, or nodes, connected to the network, no matter who controls the network overall. A digital asset payment system operating with a decentralized control model would need mutually agreed upon data governance and data privacy standards across the network. Here, a privacy strategy would only be as effective as its adoption rate throughout the network's participants, who are possibly all in some way responsible for network governance and upkeep. One possible assurance for privacy protection in a decentralized network would be a set of PETs that are deployed by default across every transaction or update within the network (here, not requiring a user to opt into them, as discussed in the next section).

4.2 Selective Disclosure

This concept is a key component of privacy-by-design. Selective disclosure encourages system architects to shift some of the decision making to the end user. The payment system provides selective disclosure opportunities, when relevant, by giving users the option to keep confidential or to disclose certain pieces of their own data to others (by selecting or not to use a PET). This empowers an individual to maintain control over their data and decide which pieces they would like to share confidentially with other actors, reveal publicly, or not disclose at all.[23] The end user, not the network administrator, could continuously opt for these selections, starting with their onboarding to acquire a digital asset and throughout their experience with a digital asset.

A digital asset design facilitating selective disclosure at the circulation layer would, as many wallet providers do, allow most privacy features to be optional. This opt-in/opt-out structure would still require the designer to select and offer a variety of privacy features. Placing this responsibility into the hands of the end user, however, also provides them the ability to disregard privacy protections to gain some other benefit, such as increasing their transaction time. A user should be made aware of the data points, such as timestamps, that may not be available for optional disclosure. Additionally, a structure that allows users to select which privacy tools they will employ and which data points they will disclose may erode the user's data privacy over time by selecting to *not* use a privacy tool, such as a shielded address, or by limiting the user's ability to participate in the system. This dichotomy illustrates the precarious nature of data privacy. The user's choices, presumably made in their own best interest, may not always be to their benefit. As a result, a system facilitating selective disclosure could assist the user in open and transparent education about how certain choices might erode or bolster a user's data privacy over time. Privacy, once eroded, is nearly impossible to reconstitute.

[23] UC Berkely has devised a set of "privacy patterns" to help organizations operationalize privacy-by-design. One such pattern outlays a comprehensive analysis of implementing selective disclosure within a system that collects both required and optional user data. For documentation, see: https://www.privacypatterns.org/patterns/Support-Selective-Disclosure

4.3 Digital Wallets

A digital wallet records a user's holdings of one or more digital assets through an internet-connected software application ("hot wallet") or a physical device not connected to the internet ("cold wallet"). A digital asset's issuance and circulation layers could possibly allow for digital wallets to assist the end user. Further, a digital asset issuer could determine if they are the sole provider of any digital wallet or if third parties could support wallets. Introducing and allowing third-party wallets into a digital asset ecosystem might decentralize the data storage and privacy burdens away from the digital asset issuer and towards wallet providers.

Digital wallets can also be custodial or non-custodial, in which either a third party manages a user's access to their funds (custodial) or only the user controls their access to their funds (non-custodial). There are benefits and risks to either custody arrangement. Custodial wallets could offer "Know Your Customer" (KYC) and auditing tools but would require more privacy considerations by the wallet issuer, such as a digital asset issuer or approved third-party, as they would likely collect personally identifying data about each wallet user. Non-custodial wallets could preserve "a direct economic relationship but not a direct technical relationship" between the digital asset issuer and users (yet still not to be confused as a guarantee of "cash-like" anonymity). But without any user identity or data collection for a non-custodial digital asset wallet, this arrangement could also facilitate untapped potential for money laundering and criminal financing.

4.4 Prevention of Illicit Activity

A system designer may have to walk a thinly balanced line between privacy protection and illicit activity prevention to avoid effectively eroding the end user's privacy for the sake of meeting jurisdictional requirements. However, intentional privacy integrations need not be a decision between jurisdictional requirements or user data protection, and anonymity need not be seen as the only acceptable privacy goal, as discussed earlier in this paper. Instead, the development of a comprehensive privacy strategy and a well-balanced selection of privacy technologies and techniques could collectively facilitate end user confidentiality while mitigating illegal activity risks.

Jurisdictions often require retail payment platforms to collect data and verify a user's real-world identity, such as the KYC process, to facilitate AML and combat the financing of terrorism. Confidentiality and jurisdictional requirements for digital assets, such as AML, are not mutually exclusive, however the initial approach to balancing requirements and privacy might be to implement a singular privacy policy meant to address a complex concern. For example, limiting "anonymous" (more likely confidential) transactions to small value payments could curtail, but likely not eliminate, the possible enabling effect

they would have for criminal activity if large value payments were allowed without identification requirements. Yet, even if a digital asset network implemented a transaction limit for anonymous or confidential payments, this would likely not be sufficient to address such a complex problem as the funding of illicit activities. In fact, an illicit actor may even self-select to break up their activity into small value transactions to hide money laundering activities, in a strategy known as "structuring."[24] Transaction value limits may have limited or unintended consequences in combating illicit activity. Instead of pursing a singular privacy-by-policy approach, then, a system designer may opt to explore a hybrid approach, in which some privacy policies are implemented alongside a robust data strategy and use case-specific privacy technologies and techniques.

4.5 Throughput and Performance

Performance goals, such as transactions per second (TPS), may alter which core privacy technologies and techniques should be chosen to help a digital asset meet its processing benchmarks. Some privacy-enhancing features may require increased processing power, which may slow transaction performance and limit potential throughput. This expenditure may be offset with a transaction fee levied against a digital asset user submitting a transaction to the system with compute-intensive privacy features. Critically, however, should the end user value, for example, a quick transaction settlement time, the user may opt out of a resource-intensive privacy feature to accelerate their transaction's settlement time. This self-selected preference towards performance over privacy may decrease the end user's settlement time while also unintentionally degrade their privacy within the ecosystem and, possibly, degrade the payment system's collective privacy. Optimistically, ongoing research into the throughput optimization of privacy-enhancing features may provide a viable path towards a compute-optimized digital asset payment system that leverages novel PETs while not requiring a stark choice between performance or privacy.[25]

5.0 Conclusion

A robust end-user data privacy strategy for a digital asset payment system starts with a thoughtful approach to technical aspects of privacy during the design phase. Anonymity, confidentiality, and full disclosure make up the spectrum of privacy, with each aspect facilitated by different design choices and confidentiality being a worthy goal system wide. A privacy strategy developed through a hybrid approach including

[24] The U.S. Treasury Department's Financial Crimes Enforcement Network (FinCEN) issued explanations of "structuring" techniques for existing financial rails as well as guidance reporting suspected structuring activities. For details, see: https://www.fincen.gov/resources/statutes-regulations/administrative-rulings/suspicious-activity-reporting-structuring

[25] This research is already underway at multiple institutions as a robust topic of discussion, in the context of both data privacy and digital asset development. Such research is also resulting in the development of new platforms, applications, and developer tools. For an example of a research effort resulting in a developer toolset, see: J. Eberhardt and S. Tai, "ZoKrates - Scalable Privacy-Preserving Off-Chain Computations," https://ieeexplore.ieee.org/document/8726497

privacy-by-design and privacy-by-policy could provide a holistic view of how, when, and where PETs and privacy tools should be employed. Specific combinations of privacy-enhancing technologies, such as fully homomorphic encryption, zero knowledge proofs, and secure multi-party computation, could provide a nuanced and novel approach to data privacy coverage across multiple tiers and use cases within a digital asset's ecosystem. Additional system technical requirements, such as throughput, or policy requirements, such as auditing compliance, may also impact a privacy strategy's needs and the resulting mix of privacy-enhancing technologies employed throughout a digital asset's system.

References

Asrow, Kaitlin, and Spiro Samonas. "Privacy Enhancing Technologies: Categories, Use Cases, and Considerations." Fintech Edge, June 1, 2021. https://www.frbsf.org/banking/publications/fintech-edge/2021/june/privacy-enhancing-technologies/.

Barker, Elaine, Miles Smid, Dennis Branstad, and Santosh Chokhani. "A Framework for Designing Cryptographic Key Management Systems." Washington, D.C.: National Institute of Standards and Technology, August 2013. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-130.pdf.

Barker, Elaine B., and Quynh H. Dang. 2015. "Recommendation for Key Management Part 3: Application-Specific Key Management Guidance," January 2015. https://doi.org/10.6028/nist.sp.800-57pt3r1.

Bindseil, Ulrich. "Working Paper Series: Tiered CBDC and the financial system." Frankfurt am Main, Germany: European Central Bank (ECB), January 2020. https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2351~c8c18bbd60.en.pdf.

Cavoukian, Dr. Ann. "Privacy by Design: 7 Foundational Principles." Ontario, Canada: Information and Privacy Commissioner of Ontario, January 2011. https://www.ipc.on.ca/wp-content/uploads/Resources/7foundationalprinciples.pdf

Cheng, Jess, Angela N Lawson, and Paul Wong. "Preconditions for a General-Purpose Central Bank Digital Asset." FEDS Notes. The Board of Governors of the Federal Reserve System, February 24, 2021. https://www.federalreserve.gov/econres/notes/feds-notes/preconditions-for-a-general-purpose-central-bank-digital-asset-20210224.html.

Davidow, Danielle, and Yacov Manevich. n.d. "Privacy-Preserving Payment System with Verifiable Local Differential Privacy." Accessed 2023. https://eprint.iacr.org/2023/126.pdf.

European Central Bank. "Eurosystem report on the public consultation on a digital euro." Frankfurt, Germany: European Central Bank, April 2021. https://www.ecb.europa.eu/pub/pdf/other/Eurosystem_report_on_the_public_consultation_on_a_digital_euro~539fa8cd8d.en.pdf#page=15, page 15.

Garfinkel, Simson L. “De-Identification of Personal Information.” NIST Technical Series Publications. National Institute of Standards and Technology, October 2015. https://nvlpubs.nist.gov/nistpubs/ir/2015/NIST.IR.8053.pdf.

Gentry, Craig. 2009. Thesis of A Fully Homomorphic Encryption Scheme. Edited by Stanford University. PDF, Stanford University. https://crypto.stanford.edu/craig/craig-thesis.pdf.

Goodell, Geoffrey, Hazem Danny Al-Nakib, and Paolo Tasca. 2021. “A Digital Currency Architecture for Privacy and Owner-Custodianship.” SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3766385.

Hastings, Marcella, Brett Hemenway, Daniel, and Steve Zdancewic. 2019. “SoK: General Purpose Compilers for Secure Multi-Party Computation.” IEEE Symposium on Security and Privacy, May. https://doi.org/10.1109/sp.2019.00028.

Johnson, Everett C., and Kenneth D. Askelson, eds. Rep. *Records Management, Integrating Privacy Using Generally Accepted Privacy Principles*. American Institute of Certified Public Accountants and the Canadian Institute of Chartered Accountants, 2009. https://us.aicpa.org/content/dam/aicpa/interestareas/informationtechnology/resources/privacy/downloadabledocuments/10252-346-records-management-pro.pdf.

Khan, Charles. “Payment Systems Privacy.” St. Louis, MO: Federal Reserve Bank of St. Louis, July 2018. https://research.stlouisfed.org/publications/review/2018/07/16/payment-systems-and-privacy.

Möser, Malte, Kyle Soska, Ethan Heilman, Kevin Lee, Henry Heffan, Shashvat Srivastava, Kyle Hogan, et al. 2018. “An Empirical Analysis of Traceability in the Monero Blockchain.” Proceedings on Privacy Enhancing Technologies 2018 (3): 143–63. https://doi.org/10.1515/popets-2018-0025.

National Institute of Standards and Technology. “Anonymity.” In *Computer Security Resource Center*. National Institute of Standards and technology, 2011. https://csrc.nist.gov/glossary/term/anonymity.

Reid, F., Harrigan, M. (2013). An Analysis of Anonymity in the Bitcoin System. In: Altshuler, Y., Elovici, Y., Cremers, A., Aharony, N., Pentland, A. (eds) Security and Privacy in Social Networks. Springer, New York, NY. https://doi.org/10.1007/978-1-4614-4139-7_10

Senate Committee on the Judiciary, John Andelin, Jean Smith, Daniel Kevin, and Fred Weingarten. Report. Edited by Sam Hale, Fred Wood, and Zalman Shaven, Selected Electronic Funds Transfer Issues: Privacy, Security, and Equity §. NTIS order #PB82-202532 (1982). https://www.princeton.edu/~ota/disk3/1982/8223/8223.PDF.

Spiekermann S., Cranor L., "Privacy Engineering". IEEE Transactions on Software Engineering, Vol. 35, Nr. 1, January/February 2009, pp. 67-82.