Hedera™

ARCHITECTING FOR PRIVACY & DATA PROTECTION ON HEDERA

AUTHOR:

Debra J. Farber

PRIVACY STRATEGIST AT HEDERA

Hedera™

CONTENTS

Hedera™

INTRODUCTION

The internet we have today is broken. Innovators of the world-wide-web developed technologies for sharing information, but, unfortunately, they did not build with privacy and security in mind. Instead, we rely on architectures based on the concept of stand-alone computers, where organizations store data centrally on a server that is sent or retrieved by a client. Whenever we interact via the internet, copies of our data get sent to the walled gardens of a server owned and operated by a service provider. We then lose control over that data.

The centralization of personal data, massive tech giants, has created a profound power imbalance. Individuals feel powerless to make meaningful choices, while Internet Giants have used behavioral advertising that tracks us, dark patterns to manipulate individuals into turning over more information about themselves, and opaque privacy notices. Furthermore, with reliance on centralized storage, we've seen the advent of massive data breaches, the rise of surveillance capitalism, and the increased distrust in centralized authorities worldwide. The EU enacted the General Data Protection Regulation (GDPR) in 2018 to address these concerns, which has become the gold standard for similar legislation. The GDPR carves out data protection rights for individuals, gives individuals more control over the collection and use of their data, and holds organizations accountable for architecting with privacy by design and default when developing products and services that process personal data.

Global researchers and developers are creating the decentralized web, or Web 3, via new protocols, including distributed ledger technology, like blockchain or hashgraph, that removes the need for intermediaries during transactions. This re-architecture of the web further democratizes its control and aims for individuals to manage their own identities and personal data. Moreover, they bring assurances that anyone anywhere can test the integrity of an asset or commodity and its current and previous ownership.

The rise of distributed ledger technology is a driving force for this movement by which developers can build decentralized applications (dApp). Although decentralized architectures introduce different application designs, the application developer still has a responsibility to understand the impacts of personal data protection.

This whitepaper serves as a guide for understanding key privacy and data protection concepts to guide you when building dApps on any distributed ledger technology network. This paper consists of three sections:

SECTION 1

Global Privacy Requirements

Understand essential privacy and data protection requirements of which architects and developers should be aware. It describes critical global data protection rights afforded to individuals and organizational obligations to respect and facilitate those rights, data minimization, accountability, security, and privacy by design and default.

SECTION 2

Architectural Considerations

This section describes emerging architectural patterns for deploying decentralized applications and three privacy and data protection pitfalls to avoid: 1) who is accountable for compliance, 2) deletion of personal data from the blockchain or hashgraph, and 3) the finality of smart contracts made in error.

SECTION 3

Inheriting Trust via Hedera

Hedera offers a variety of network services to interact with the public ledger. For each network service, there are relevant privacy considerations based on how the data flows through the network and its deployment options.

SECTION 1

GLOBAL PRIVACY REQUIREMENTS

Global data protection laws like the EU's General Data Protection Regulation (GDPR), Brazil's Global Data Protection Regulation (LGPD), California's Consumer Privacy Act (CCPA), India's Personal Data Protection Bill (PDPB), China's Personal Information Protection Law (PIPL), and other emerging regulations largely follow the same privacy principles across jurisdictions. These regulations require organizations to design and develop products and services with a "data protection-by-design and default" approach (AKA "privacy by design"), outline new rights afforded to individuals to which organizations must respect, and include accountability requirements that require organizations to demonstrate their compliance.

Privacy by Design and Default

When adopting a public ledger, architects and developers have taken two distinct approaches. More often seen in enterprise deployments, the first approach leverages a distributed ledger to increase the transparency or trust of a currently centralized process. The second approach, which is viewed as more decentralized, leverages an entirely decentralized stack that gives control directly to individuals instead of centralized organizations. While this paper focuses on the first approach, as it enables enterprises to adopt distributed ledger technology more quickly, every developer should be aware of their obligation to build products and services as legally required by GDPR and most international data protection laws.

When building on a public ledger, you should follow the 7 Privacy by Design Principles for compliance with the principles that underlie privacy and data protection laws.

1. PROACTIVE PRIVACY ARCHITECTURE, NOT REACTIVE

You must understand the privacy impacts on the product or service you develop: the threat model for potential privacy harms, mitigate your privacy risks, plan out how to protect and secure personal data that you process, and keep up-to-date with privacy-enhancing strategies, tactics, and deployable technologies. Like security, you cannot successfully bolt on privacy after you build a product or service. You are obligated to prevent privacy risks from occurring in the first place.

2. PRIVACY BY DEFAULT

To have privacy by default means that you require no action to enable higher privacy levels because defaults provide the most significant privacy level of control. For instance, let's imagine you are building a social sharing feature with three settings - share publicly, share with friends, or share only with me. You are obligated to set the default to share only with me, allowing users to choose a lower level of privacy as they see fit.

3. PRIVACY EMBEDDED INTO IT SYSTEMS AND BUSINESS PRACTICES

It is a mistake to attempt to bolt-on privacy to an already-built system. It usually never works. Instead, turn privacy into an essential component of your dApp's core functionality on which you deliver. Make privacy integral to your dApp without diminishing functionality. Moreover, you should demonstrably minimize any privacy impacts of the developed technology, operation, or information architecture and their uses and ensure that privacy is not easily degraded through use, misconfiguration, or error.

An all too common example in the industry is sending cryptocurrency to the wrong wallet address. This misconfiguration causes actual harm to crypto speculators today, and the industry should seek to remedy this privacy vulnerability. As an architect, you should review this privacy problem and consider strategies, tactics, and techniques to reduce human errors (e.g., via automation, etc.).

4. FULL FUNCTIONALITY

Privacy by design seeks to accommodate all legitimate interests and objectives in a positive-sum "win-win" manner, not through the dated, zero-sum (either /or) approach, where we force unnecessary trade-offs. If you sacrifice functionality for privacy, then you are doing it wrong. Make sure to avoid the pretense of false dichotomies, such as privacy vs. security. Instead, demonstrate that it is indeed possible to have both.

5. END-TO-END SECURITY THROUGHOUT THE DATA LIFECYCLE

This principle underscores the importance of maintaining adequate security controls for the enforcement of privacy policies. Remember that your organization assumes responsibility for the security of any personal data throughout its entire lifecycle, consistent with standards, including secure destruction methods, appropriate encryption, and robust access control and logging methods.

6. VISIBILITY AND TRANSPARENCY

All stakeholders need assurance that your dApp and supported technologies and processes operate according to the stated promises and objectives, subject to independent verification. The Visibility and Transparency principle focuses on three foundational areas: accountability, openness, and compliance

Accountability: Your company remains responsible for all privacy-related policies and procedures, which must be documented and communicated as appropriate, and must be assigned to a specified individual's oversight. Further, you ensure that individuals are made fully aware of the personal data collected and for what purposes, and those operations remain visible and transparent. When transferring personal data to third parties, your organization must demand an equivalent level of privacy protection via contractual clauses.

Openness: Companies must make information about the policies and practices relating to managing personal data readily available. This includes your dApp privacy policy, just-in-time notices, and maintaining accurate and up-to-date records of processing activities.

Compliance: Per the GDPR and other global data protection laws, organizations are required to put in place processes and mechanisms for handling individual complaints and redress (i.e., a dispute mechanism) from your customers.

7. RESPECT FOR USER PRIVACY

Respect for User Privacy extends to the need for human-machine interfaces to be human-centered, user-centric, and user-friendly so that informed privacy decisions may be reliably exercised. Similarly, your business operations and physical architectures should also demonstrate the same degree of consideration for individuals, who should feature prominently at the center of operations involving collections of personal data. dApp architects and operators should keep the interests of individuals paramount by offering substantial privacy defaults, appropriate notice, and empowering user-friendly options.

Understanding Data Protection Requirements

explains what personal data categories, if any, your company collects via the dApp and, under GDPR, what lawful basis permits you to process this information.

There are six lawful bases for which you may process personal data under GDPR and similar data protection laws:

1. CONSENT
2. CONTRACT
3. LEGAL OBLIGATION
4. VITAL INTERESTS
5. PUBLIC TASK
6. LEGITIMATE INTERESTS

Without one of those lawful bases, you are forbidden from processing that data. This differs from laws in countries like those in the U.S.; however, as more countries adopt the GDPR framework from the EU, this is a crucial requirement to keep in mind.

In addition, you must include a description of a person's data protection rights in your privacy notice and how they can reach out to your company to exercise those rights. Make sure to work with local attorneys to draft this notice for each jurisdictional area where you decide to place your dApp into the stream of commerce. For instance, a GDPR-focused privacy notice will not be sufficient in the U.S., and a California privacy notice will not be adequate in the EU. A privacy and data protection attorney can also guide you on the appropriate legal strategies to take depending on whether the law classifies your company as a "data processor" or a "data controller." Your organization is responsible for: the personal data that it processes; what data gets sent to the blockchain or hashgraph; respecting personal data protection rights; maintaining records and privacy and security assurances; and embedding privacy into the design of your dApp and associated business processes.

Obligation to Facilitate Data Protection Rights

The GDPR and similar global privacy laws grant rights to individuals that give them more control over how their personal data is used, and dApp companies must build data protection rights fulfillment into their technology and processes, supported by proper individuals.

Those rights include:

1. The **right of access** is when someone asks you for a copy of the data you have on them. This is also known as a data subject access request or DSAR.

2. The **right to object** means people can object to specific processing of their personal data, so you'd have to stop using their data for particular purposes unless you have a good reason to continue.

3. The **right to be informed** usually means that you have to tell people that you have their data and what you're doing with it.

4. The **right to rectification** means people can ask you to correct their data if it isn't accurate.

5. The **right to erasure** (i.e., deletion) permits individuals to request deletion of their personal data. It is also known as the 'right to be forgotten' and means that you may have to delete their data upon request in certain specific situations.

6. The **right to restrict processing** means that you have to temporarily stop processing someone's data if they ask you to. You can store their data but not use it. This isn't an absolute right and only applies in certain circumstances.

7. The **right to data portability** gives people more control over their data and where it's stored. It's intended to make it easy for them to provide it to another data controller if they need to. The data you hold about them electronically has to be made easily accessible and transferable.

If personal data is processed entirely by automatic means and this might have a legal or similarly significant effect on the person, they can request some human involvement in the processing.

Under California's CCPA, California consumers also have the right to restrict the sale of their personal information to third parties. Thus, companies who sell personal data to third parties must include an opt-out link to this sale on their main website and respect those opt-outs.

California's CCPA also grants a right to non-discrimination for exercising one's CCPA rights.

SECTION 2

ARCHITECTURAL CONSIDERATIONS FOR BUILDING PRIVACY INTO DAPPS

Traditional databases store data in one place, where security systems control access. Public ledgers store data in many places for users to control. Those who develop on public ledgers must consider the tradeoffs and advantages of the system to design a solution that best fits their use case requirements. With the transition from web 2, defined by centrally controlled platforms and walled gardens to web 3, there are a few early, yet emerging architectural patterns. The good news is that these patterns go hand-in-hand with designing privacy-centric applications that protect your users.

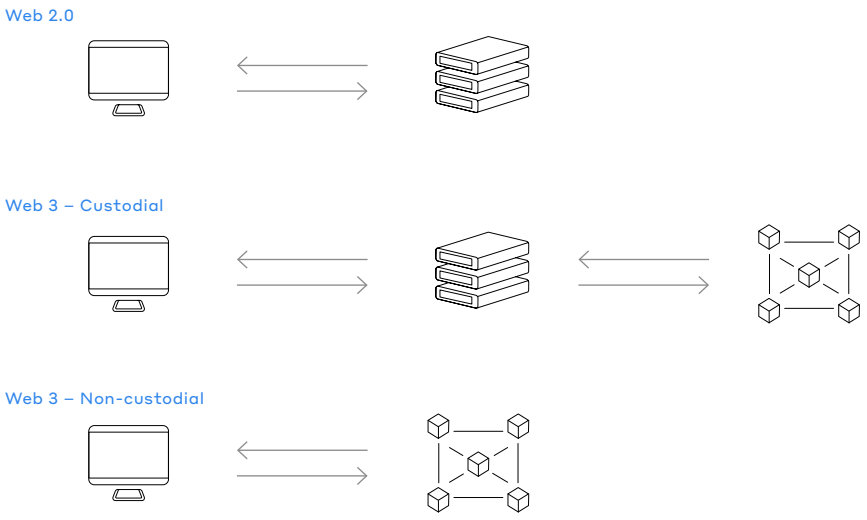

Web 3 — Custodial

A middle ground is a web 2-like authentication model that ties an individual's information, typically an email, to a set of keys and the centralized party is responsible for their design or safekeeping. Ideally, the centralized party that controls the application does not store these keys; instead, the keys are encrypted and require a person to log in to sign transactions. Some applications, such as centralized exchanges, do store a person's keys and therefore open up greater liability to a breach in comparison with its non-custodial counterpart.

Web 3 — Non-custodial

The most decentralized way for a user to interact with an application in web 3 is to have a peer-to-peer design. This design often relies on components such as a user-controlled browser extension to hold the user's keys and sign transactions and a smart contract to perform application logic. In this model, everything to conduct a transaction is stored entirely on-ledger. The user does not interact with a centralized entity but directly with the decentralized network or smart contract.

From a privacy standpoint, these types of applications are ideal. The user is entirely pseudonymous throughout the interaction, and no personal information is exchanged in the process.

Minimize Personal Data for Greater Privacy

No matter your choice of architecture, prioritizing privacy is essential. To reduce risks to individual privacy, data protection laws require organizations to follow data minimization strategies, ensuring that the personal data processed is *adequate* to properly fulfill your stated purpose, *relevant*, in that it has a rational link to that purpose, and *limited* to what is necessary for your stated goal. Further, the period for which the personal data are stored must be limited to a strict minimum. This makes sense, considering that the more personal data that is collected (or derived) about an individual and used, the more obligations your organization has to protect and secure that data. Make sure to understand your data retention responsibilities as well, since it's up to the data controller to delete data when it is no longer useful. You are required to minimize the amount of personal data collected or derived from individuals, whether or not you intend to store it on the blockchain or hashgraph.

Wherever possible, minimize the identifiability, observability, and linkability of personal data. There are several ways to achieve data minimization with distributed ledger technology:

COLLECT LESS DATA

Only collect personal data that is essential for the use of your dApp or business. By reducing the amount of regulated data you collect upfront, you lower the regulatory and privacy compliance burdens needed to protect and secure that data.

ENCRYPTION DOES NOT MINIMIZE DATA

As we should know, the encryption of personal data enhances the security of that data should it fall into the hands of an unauthorized user. However, when it comes to minimizing personal data following global data protection laws like GDPR, simply encrypting personal data does not minimize the amount of data you have (despite minimizing risks to confidentiality). The U.S. privacy laws view the personal information that undergoes encryption as anonymous data - i.e., a safe harbor in the event of a breach. However, the EU considers the cyphertext here as encrypted personal data because it's still identifiable and covered by the GDPR. This is important to keep in mind as you architect dApps for a global audience across many jurisdictions.

LEVERAGE TOKENIZATION

One benefit of tokenization is the ability to remove sensitive payment or personal data from business systems. For example, tokenization in banking protects cardholder data. When processing payment using the token stored in your systems, only the original credit card tokenization system can swap the token with the corresponding primary account number (PAN) and send it to the payment processor for authorization. These systems never record, transmit, or store the PAN, only the token. So, if a token is lost outside of the tokenized system, no data breach will result, and a new token can be issued.

Rather than storing personal data on the ledger directly, consider storing tokenized data instead. Remember that while the immutability of a distributed ledger is a feature, the immutability of personal data is incompatible with global data protection laws due to the right to erasure and rectification. Therefore, many deployments leverage distributed ledgers as an access-control moderator to personal data and use an off-chain solution to store private and sensitive data.

USE PSEUDONYMIZATION

Pseudonymization is a great minimization technique that replaces apparent identifiers, such as name or email address, with a simple reference number. Thus, while individuals are not identifiable from the dataset itself, they can be identified by referring to other information held separately. Pseudonymous data is therefore still considered personal data, and data protection laws apply.

An overarching benefit to pseudonymization as a data minimization technique is that it simplifies your data protection compliance. Pseudonymization techniques can reduce the risk of harm to individuals that may arise from personal data breaches, which in turn reduces your company's breach notification burden. In addition, pseudonymization may narrow the sheer volume of data that must be considered when responding to access and deletion requests from individuals. Typically, dApp developers use cryptographic hashes to appropriately pseudonymize on public distributed ledgers. You can use hashing to create a unique one-way hash of arbitrary digital data that lives on the public ledger. While the hash is uniquely tied to an identity that remains off ledger and can easily be computed given that data, the fingerprint by itself cannot be used by anyone to recreate the data to which it corresponds. Hashes enable dApp developers with flexibility to compute data tied to an identity while preserving privacy of that identity on public ledgers.

ANONYMIZE YOUR DATASETS WHEN SHARING OUTSIDE YOUR ORGANIZATION

Data protection law does not apply to anonymous datasets since they no longer identify individuals, thus eliminating any risks to their rights and freedoms. Therefore, many organizations choose to anonymize personal data; you must use techniques that reduce the chances of identifying individuals to a sufficiently remote level where they cannot be re-identified. Anonymization limits data protection risks and can enable you to safely make information available to organizations outside your own or even to the general public without violating any privacy rights. It is generally easier to disclose anonymous information than personal data as fewer legal restrictions apply. You can also use anonymous information in new and different ways, as the data protection rules on purpose limitation do not apply to anonymous data. Use cases where anonymization of data sets is appropriate or required include: analytics, research, sale to third-parties, or sharing data publicly.

Accountability: Demonstrate Privacy Compliance

The GDPR and similar laws require data controllers to adhere to the Accountability principle, which requires that organizations *demonstrate* their compliance with the suitable control measures and records. "Privacy by design and default" has long been a best practice when designing new products, services, and systems that process personal data. The GDPR and similar legislation have now enshrined it into law. Organizations must follow this design approach because data protection by design and default is essential for accountability. The integration of data protection considerations into your operations will facilitate compliance with your obligations while transparently documenting the decisions you take. Privacy and security teams can assess data protection risks and suggest appropriate measures to mitigate those risks, such as minimizing the data you collect, applying pseudonymization techniques, leveraging zero-knowledge proof strategies, and improving security features.

RECORDKEEPING REQUIREMENTS

Global privacy and data protection regulations have numerous recordkeeping requirements, which include the need to: maintain an up-to-date privacy notice that is easily accessible; adopt and implement data protection policies; put written contracts in place with organizations that process personal data on your behalf; conduct impact assessments to determine risks to privacy; and maintain documentation of your organization's personal data processing activities.

SECURITY

Your organization is responsible for putting in place *appropriate* security measures to prevent the personal data that you hold from being accidentally or deliberately compromised. To determine whether your security measures are reasonable, you need to take a risk-based approach. You should review the personal data that you hold and how your organization uses it to assess how valuable, sensitive, or confidential it is and the harms it may cause if that data was compromised.

PRIVACY ASSURANCES

Privacy assurance is the measure of confidence that the data protection features, practices, procedures, and architecture of an information system accurately enforce the privacy policy. Assurance is determined by the evidence produced by your assessment process. Evidence can be stored in logs, as a hash on the blockchain or hashgraph, in databases, or elsewhere. Thus, make sure that you can demonstrate the effectiveness of your privacy and data protection controls by making evidence easily accessible and available. For example, record when a customer consented to collecting and using their data for a particular purpose or requested deletion of their data.

SECTION 3

DATA ON HEDERA

When it comes to decentralized networks, it's paramount to consider the technology's capabilities and its governance structure to understand how it is likely to evolve to meet changing regulatory requirements. This section provides an overview of how Hedera and its network services manage data to help you understand when and how to think about protecting privacy and compliance with data protection requirements when deploying your applications.

HEDERA ACCOUNTS

The basis for a user or application interacting with Hedera is a Hedera account. A Hedera account number serves as a pseudonymous identifier, such as 0.0.1234. This pseudonymous account ID lives on the ledger and is controlled by asymmetric cryptographic keys. This model, which is shared across all blockchains and public ledgers, encourages the option to remove the typical Web 2.0 interaction of an email, password, and other associated personal metadata required to interact with an application. Thus, Hedera accounts limit the amount of data stored by a centralized party to the bare minimum.

DATA PERSISTENCE AND NODE TYPES

Hedera networks consist of two types of nodes: consensus nodes and mirror nodes. Consensus nodes receive transactions from clients, charge transaction fees, and contribute to the achievement of consensus. Consensus nodes are responsible for preserving the network's most critical information – user accounts, their hbar and crypto balances, smart contracts, and files. They were designed to reduce storage bloat and ensure predictable performance as network utility continues to rise.

On the other hand, mirror nodes are read-only nodes that receive information directly from consensus nodes. They offer the ability for the broader Hedera ecosystem and individual application owners to store the entire history of the ledger, for instance, to track account balances throughout time or the full provenance of a supply chain.

Hedera services

Organizations that build on Hedera often maintain their services, databases, and privacy policies, and are accountable for incorporating privacy requirements into their processes, design, and architecture. It is up to the developers, companies, and individuals who build on Hedera to be responsible for architecting their services with privacy and trust in mind.

Through its network services and supporting functionality, Hedera does provide more significant opportunities for developers to choose what's suitable for their customers and use cases.

HEDERA CONSENSUS SERVICE

Hedera Consensus Service accounts for the majority of transactions on Hedera today. It offers the ability to define a topic and send messages, arbitrary event data, to the topic.

Developers can encrypt messages sent to the hashgraph as a much-needed control that enforces data minimization and the confidentiality of personal data. Note that while a developer can encrypt a message containing personal data to meet some global data protection requirements, their organization is still required to demonstrate their accountability for that encrypted data - i.e., unlike in the U.S., encrypted personal data is still considered personal data in the EU and elsewhere.

Any personal data you send via Hedera messages will persist on the mirror nodes, which dApp developers use to query historical transactions. Therefore, application owners must understand the data protection problems when including personal data in Hedera messages and act accordingly. Most definitions of personal data include any linked or linkable data to an identifiable individual, which is quite broad and can include everything from an IP address, specific metadata, or even one's shoe size. Make sure to get advice from data protection counsel and operational privacy experts to understand what data you should keep off-ledger and how to architect your dApp for global data protection compliance.

HEDERA TOKEN SERVICE

Hedera Token Service (HTS) enables the configuration and management of fungible and non-fungible tokens (NFT). For tokens on Hedera, like every transaction, they're transparently verifiable. This means that all ecosystem participants can freely see information about each token and its transfers.

By their very nature, NFTs are pseudonymous since NFTs are linked to the owner's account ID. To minimize risks to privacy exposure, global data protection regulations require that companies only collect and process the minimum personal data necessary. Beyond art and real estate, organizations can use HTS to mint NFTs that represent and provide authenticity for any real-world asset, including a data protection right or a digital right, and bind them to pseudonymous digital identities to reduce the number of instances where employees store personal data across the enterprise. Effectively, the advent of NFTs makes it possible for companies or the public to verify that a person or entity meets certain criteria without actually sharing specific information that could be vulnerable to being compromised.

Balaji Srinivasan refers to the growing NFT ecosystem as the "Pseudonymous Economy." As we continue to associate NFTs with our online identities, which we manage ourselves, we move from interacting online with our identity cards to our pseudonymous accounts. It is important to remember that pseudonymization of data is an essential technique for minimizing the attack surface of personal data. However, pseudonymous data is still considered personal data because it is still tied back to an identity. Only anonymous data falls outside of the GDPR and most global regulations. So, companies that process personal data and add reference hashes on a ledger must still meet their many other privacy and data protection obligations.

Hedera designed HTS based on user requirements from existing token issuers and interested enterprises, which included greater scalability at a lower cost and several compliance-focused features.

HTS allows each token type to have several options for issuers to choose from several levels of control. These range from being entirely immutable and decentralized to having the flexibility to mark accounts for KYC verification, freeze their token holdings, and manage token supply, among others. In the instance of the KYC parameter, if a company decides to issue a custom token and implements KYC verifications where a third party verifies an account owner's identity, then the account will update a flag on the account on a per-token type basis to mark it as KYC'd. If KYC is required, then a not-KYC'd account cannot send or receive the token. The KYC key must sign the KYC transaction.

HEDERA SMART CONTRACT SERVICE AND FILE SERVICE

Data protection regulations require companies that utilize any smart contracts to communicate and correct any errors. Since smart contracts and files on Hedera can be permanent or have binding arbitration, you may need to nullify and replace a smart contract or consider a governance structure to control a set of multi-sig keys to modify the smart contract. Make sure your organization has the appropriate rights to nullify & replace smart contracts or files as needed.

CLOSING

To support Hedera's developer community, Hedera remains committed to helping developers thrive by providing resources, tools, and features that support developers' privacy-by-design requirements and data protection accountability.

Furthermore, throughout this paper, we've focused on many of the technological components of Hedera but must acknowledge the importance that governance plays in a decentralized network. A general-purpose public ledger should be governed by representatives from a broad range of market and geographic sectors, each with world-class expertise. Those responsible for network governance need technical knowledge so they can competently manage the platform's underlying software. They need business and economics expertise so they can drive business operations of the ecosystem. They need legal expertise to help navigate the evolving regulatory environment. In other words, the network should be governed by a decentralized group of globally recognized industry leaders, representative of every market in the world. As part of this, the Hedera Governing Council is well-positioned to prioritize the privacy needs required to meet global regulations best today and in the future.

Public ledgers are modern technology enabling architects and dApp developers to build secure, privacy-preserving, and compliant systems. That said, they must still recognize the data protection guardrails to put in place and carefully architect their processes that use personal data so as not to cause: data breaches; misuse of personal data for purposes beyond their allowed-for use; the inability to delete personal data from the hashgraph or blockchain; or other privacy and data protection problems that are difficult to change later - i.e., avoid creating risky privacy technical debt.

For enterprises endeavoring to build on distributed ledger technology, it is essential to build out privacy engineering functions that educate product developers and engineers on how to implement privacy by design into your tech stack from the bottom layer up. Once enterprises appropriately evaluate their threat model for privacy and data protection risks, put in place controls and supporting processes, and design for privacy and data protection by design and default, then they will demonstrate to their customers, employees, regulators, legislators, and the public at large that they care about privacy and aren't just playing a whack-a-mole game of compliance for the sake of compliance.

GLOSSARY

ANONYMOUS DATA	Anonymous Data is data unrelated to an identified or identifiable individual (i.e., data that is not "personal data").
CENTRAL BANK DIGITAL CURRENCY (CBDC)	A central bank digital currency uses an electronic record or digital token to represent the virtual form of a fiat currency of a particular nation or region. A CBDC is centralized; it is issued and regulated by the competent monetary authority of the country.
DATA CONTROLLER	The Data Controller is the person or organization that exercises overall control over the *purposes and means* of the processing of personal data.
DATA PRIVACY	Data privacy generally means the ability of a person to determine for themselves when, how, and to what extent personal information about them is shared with or communicated to others.
DATA PROCESSOR	The Data Processor does not have any purpose for processing personal data and only acts on a client's instructions.
DATA PROTECTION	Data protection is a set of strategies and processes you can use to secure your data's privacy, availability, and integrity. In countries like the EU, the concept of privacy is enshrined as a fundamental human right protected by the General Data Protection Regulation.
HEDERA CONSENSUS SERVICE TOPIC	An Hedera Consensus Service topic manages the stream of messages for one application, such as a market where people bid on products.
NODES	Nodes are decentralized computers that store a copy of the distributed ledger.
PERSONAL DATA	Personal data is about who you are, where you live, what you do, and more. It's anything linked or linkable to an identifier.
PSEUDONYMOUS DATA	Means personal data that cannot be attributed to a specific data subject without the use of additional information, as long as such additional information is kept separately and subject to technical and organizational measures to ensure non-attribution.
TOKENIZATION	Tokenization is the process of removing sensitive data from business systems by replacing it with an undecipherable token and storing the original data in a secure data vault.