How Fintech Companies Can Break the False Dichotomy of Data Privacy Versus Data Utility

skyflow

Contents

Abstract

Authored by Harsh Karmarkar, Head Of Solutions, Skyflow

Modern fintech companies face ever-increasing complexity — from hybrid data infrastructures to a growing list of compliance requirements. They have a vital need to protect sensitive customer data while not losing data utility.

At the same time, fintech companies sit at the intersection of new technology and the constant reinvention of financial services. The need for rapid fintech innovation has never been greater, especially given recent global trends like:

- An accelerated shift to digital and remote work, hastened by the COVID-19 pandemic
- The demand for financial services to meet the needs of unbanked and under-banked customers
- New competition and business models coming from disruptive Web3 companies
- Rapidly rising inflation and interest rates, which create demand and increasing viability for new credit offerings

With this rapid innovation comes heightened potential for risk. The customer interaction, transaction, and business process insights that fintech companies need to operate and innovate are dependent on their customer's most sensitive financial data. Meanwhile, this is the same data that's sought by malicious hackers.

Recent data breaches have made consumer trust a competitive battleground. One recent study found that nearly 40% of customers reported losing trust in a company because of reported data breaches. That study also found that 88% of customers are willing to walk away from a company that they don't trust.

The tension between using data to innovate and the risks of using sensitive data has created a false dichotomy between the utility of sensitive data and the need to maintain the privacy and trust of the customers to whom it belongs. Why is this a false dichotomy? Because recent innovations make possible a new and more sophisticated approach to data privacy and protection that doesn't sacrifice data utility.

Forward-looking fintech companies are leveraging new architectural patterns to break this false dichotomy by isolating and protecting sensitive data in a data privacy vault. By centralizing sensitive data in a vault, they can use tokenization, encryption, and granular data governance to meet this complexity head-on. This allows them to maximize the usefulness of sensitive data without putting its privacy at risk.

The Origins of the Data Security Versus Data Utility Paradigm

The false choice between data protection and data utility comes from a very real historical context. When data went from being something physically recorded on paper to being stored digitally, there were a limited set of tools available or needed to secure it. Information security was still governed by physical access to machines and media. Security became even more difficult when data became ‘networked’ — shared across machines within a network or with the world as a whole via the internet. The most common perspective on security was to view it as a binary choice: either data was secured or it was not. Given that the world of physical security often revolves around binaries like “is the door locked?”, this isn’t too surprising.

For a long time, the only tools available to security and IT professionals were encryption and access control. Most internal systems would have a small set of users who could decrypt encrypted data, while everyone else was denied access. This was the standard approach that still leveraged the binary view of data security, in which a given dataset is fully accessible to some users, completely inaccessible to other users, and a point of frustration for those who only needed partial access. The limitations are obvious: the flow of data is bottlenecked by a few privileged users, limiting its utility. And worse yet, anyone who could gain access to those privileged user accounts would have unlimited access to sensitive data.

Tokenization

Tokenization offered a different approach — rather than encryption, which uses an algorithm to transform the data into something that isn’t legible, true tokenization independently generates a token that serves as a “stand-in” that maps back to the original data. Tokens are independent, so you can’t “brute force” them. You can also preserve the format of the original data, so a tokenized email address looks like an email address but is otherwise unreadable, allowing you to add different increments of differential privacy. For example, you could fully tokenize an email address by tokenizing both the username and domains, or you could just tokenize usernames so that the domain remains readable from the token.

Most importantly, tokenization allows you to replace sensitive data at a field level in systems that are designed to expect plain text data (like an email address) in a specific format. Once it became available, security experts could pair tokenization

with encryption to greatly increase data utility and remove some of the bottlenecks to data access. To learn more about tokenization, see Demystifying Tokenization: What Every Engineer Should Know.

RBAC and Redaction

More sophisticated tooling and a more nuanced understanding of data-intensive workflows allowed security architects to discard the binary view of data security in favor of a more nuanced view of the utility of each element of a given sensitive data field. They learned that protection can be incremental: for example, each part of a payment card number can be made accessible to different users for different use cases without making it either fully unprotected, or fully inaccessible. The key to this is a combination of governance (or access control) and redaction (sometimes called masking). Adding governance and redaction to tokenization means that security isn't an all or nothing binary any longer.

Say, for example, your company employs a customer service agent (CSA) who needs to choose between two stored credit cards in a customer's profile. It probably isn't necessary to decrypt the entire credit card number — they could verify it using just the last four digits. In this context, this verification workflow requires that the CSA can view a subset of this sensitive data. You can eliminate a bottleneck by giving them access to the last four digits while protecting the data privacy of the full credit card number.

The concept of data utility is that the value of data varies with its context. Knowledge of an individual's intent to refinance their home, for example, is worth much more to a lender than it is to a burger chain. The utility of the data is based on the value of the decisions or operations it enables, and for whom.

This nuanced approach, combined with sophisticated technologies like tokenization and polymorphic encryption, has allowed security professionals to transcend the binary all-or-nothing paradigm that once characterized problems around data protection. Expanding the paradigm of data privacy and protection is where the data privacy vault comes in.

Enter the Data Privacy Vault

A modern security posture enables you to discard the security versus utility dichotomy by combining a sophisticated set of security tools with a centralized and secure data privacy vault. A data privacy vault allows you to isolate, protect, govern, manage, monitor, and use sensitive data for critical workflows.

Isolate

Isolating raw sensitive data away from your systems might seem like it would reduce its utility, but it is actually an essential step towards promoting data usability. Centralizing and isolating sensitive data in one place gives your team invaluable insight into how data is collected, who is using it, and how it is being used — all crucial pieces of information when governing data use.

Centralizing data also allows you to avoid one of the major issues with data security: sensitive data sprawl. Data sprawl occurs when sensitive data is replicated from one system to another, increasing the amount of infrastructure that's impacted by regulatory compliance and increasing the attack surface area for malicious hackers to exploit. At its worst, data sprawl allows any user or service with access to any part of your infrastructure to access sensitive data.

Secure

Polymorphic encryption, tokenization, and redaction may not on their own break down the usability versus security dichotomy, but they can be combined to secure data while still maintaining its usefulness. Rather than simply locking the data down, these tools provide differential privacy, simultaneously securing data while allowing it to be used.

Govern

With sensitive data isolated in a data privacy vault, access is controlled using a combination of zero trust architecture and role-based and account-based access controls (RBAC and ABAC). These access controls ensure that only the minimum amount of data that's required for business-critical workflows are available to your users and services.

What You Need to Build a Data Privacy Vault

Building a full-featured data privacy vault is a demanding task, requiring time and expertise to design a privacy-preserving data architecture and operational controls.

Expertise

While the proliferation of security tools has taken security well beyond a basic approach to encryption and access management, developing and integrating these tools without disrupting ongoing business requires a great deal of expertise. Building a data privacy vault requires significant expertise and time.

Features like data governance, encryption, and tokenization need to be implemented and administered with great care to ensure both data security and data utility. And the stakes when building a data privacy vault are unusually high: when handling sensitive data, there's always the risk of a data breach to consider. And also the risk of losing the use of internal systems that depend on encrypted data.

Architecture

A data privacy vault centralizes sensitive data, solving part of the problem. A well-designed data privacy vault also gives you differential privacy. When you're building a data privacy vault from the ground up, you're building it with the aim of preventing any sensitive data leaks, or replicating sensitive data in many places (sometimes called "data sprawl"). Beyond the basics like secured storage and an access layer, you need:

- A flexible and secure tokenization processes that lets you protect sensitive data without breaking your internal systems or creating a bottleneck behind a specific set of users
- Sophisticated, fine-grained access controls provided by a governance engine that follows zero trust principles, granting the bare minimum of data necessary to team members and giving your security team granular control
- The ability to integrate your vault with trusted 3rd party services, allowing them to utilize sensitive data without you sacrificing control or increasing security risk

What You Need to Build a Data Privacy Vault

Continued

Operations

A well-designed data privacy vault provides logging, monitoring, and fine-grained data governance - all of which requires a lot of planning to implement well. While some elements of these processes can be automated, you will still need to draw upon the combined expertise of your security team to proactively analyze, invest, and change the way you are protecting sensitive data as new products and systems are built. They also need to be able to invest in threat detection and incident response, shifting left as much as possible in order to head off any potential security or compliance issues.

Incorporating a Data Privacy Vault Into Your Architecture

To integrate a data privacy vault into existing architecture, it is best to start by choosing a small set of use cases. You'll have to extract data from a few discrete systems and replace it with tokens, giving you direct insight into what data those systems actually need, and more importantly, what can be redacted or removed from their scope of access.

Integrating a data privacy vault also requires internal education to inform employees that sensitive data access is based on what's strictly required, giving them a process to request access to the data that they need, and helping employees to understand that their workflows don't require unlimited access to sensitive data.

Education is also critical to illustrate to your team members that restricting unnecessary access to sensitive data actually adds to flexibility. By protecting sensitive data with a zero trust model and fine-grained data governance, you make it easier to embrace remote work without increasing data security risks.

Give Skyflow a Try

As a modern fintech company, you feel the simultaneous pressures of tough compliance and security requirements and the need to innovate rapidly to compete. The data privacy vault is the ideal tool to satisfy the former while enabling the latter — it's an architectural pattern that will empower your security team to protect sensitive data while also giving the rest of your company the data they need to do their jobs and build great products.

Despite these substantial benefits, building a fully-functioning data privacy vault is a very heavy lift. And building one yourself can come at the cost of resources better used to innovate and improve your core products and services. This is where Skyflow comes in.

Skyflow's best-in-class Fintech Data Privacy Vault allows you to ensure the privacy and security of sensitive customer data without sacrificing data utility. Whether you are just starting out or you need to integrate Skyflow with your already extensive architecture, our API-first approach makes it easy and allows you to see benefits right away.

If you're interested in learning how we can help you break the false dichotomy between data privacy and data utility, contact us.

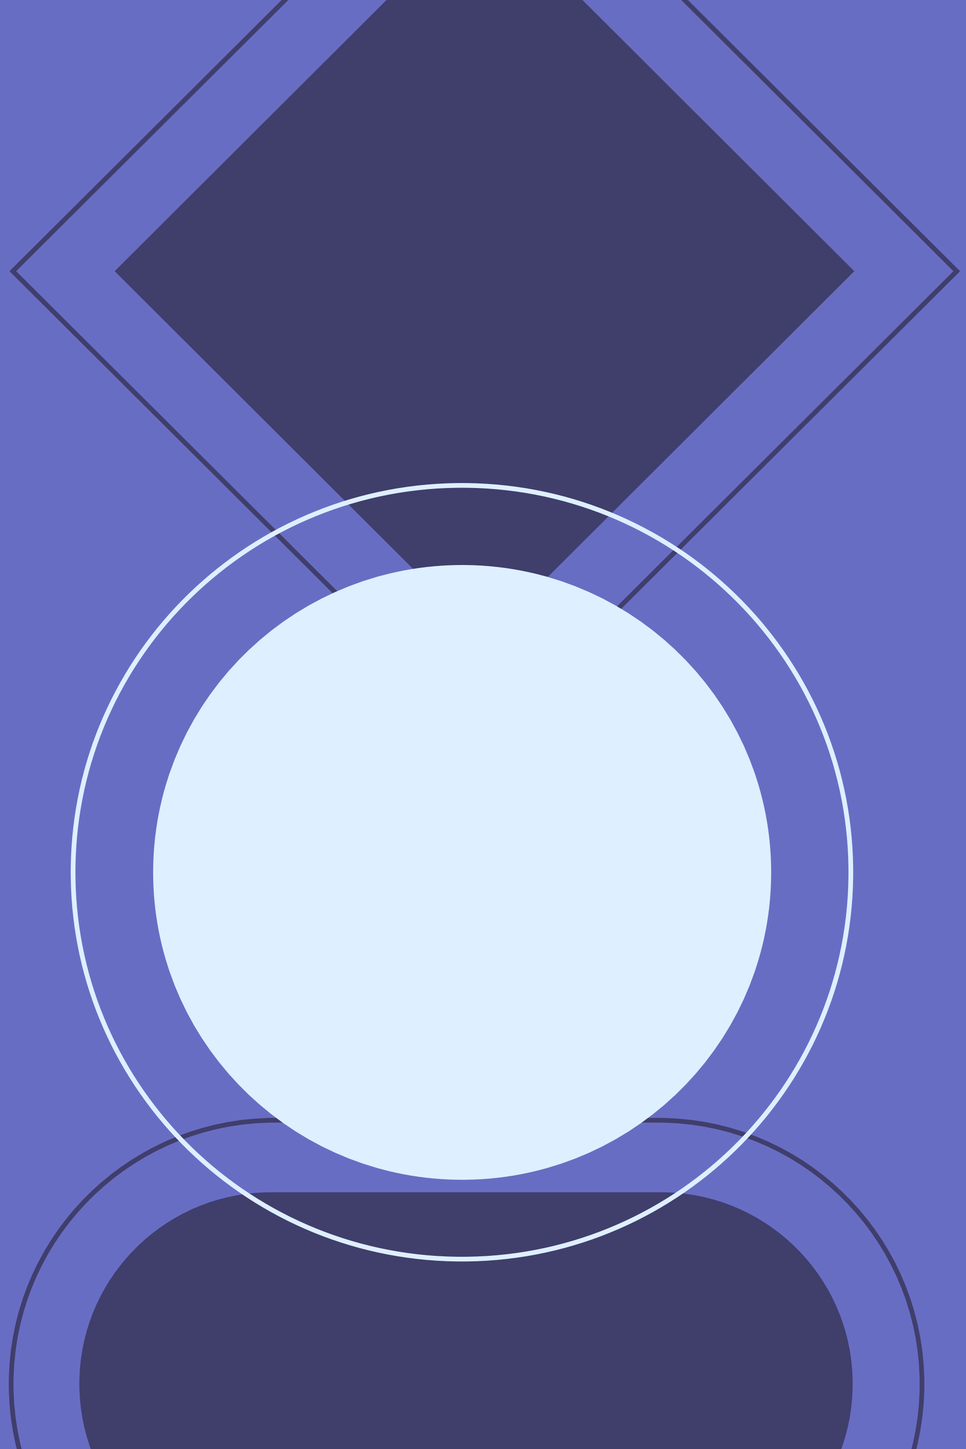

About Skyflow

Founded in 2019, Skyflow is a data privacy vault for sensitive data. The company was founded by former Salesforce executives Anshu Sharma and Prakash Khot to radically transform how businesses handle users' financial, healthcare, and other personal data that powers the digital economy. Skyflow is based in Palo Alto, California, with offices in Bangalore, India. For more information, visit skyflow.com or follow on Twitter and LinkedIn.

About the Author

Harsh Karmarkar is an architect and technical sales leader whose work has spanned multiple industry sectors including healthcare, financial services, retail, tech, media, and telecom, across North America, Asia, and Europe. He currently leads the Solutions team at Skyflow, helping organizations understand how to protect their customers' sensitive data.

skyflow