



KARLSGATE IDENTITY EXCHANGE

Technical Whitepaper

Enable agile data collaboration at scale, seamlessly connect data from diverse sources and power AI learning engines with advanced automation for precision and speed.

Introduction

Privacy and information security are emerging as the next big challenges of the information age. For decades, data has been shared, posted, and transmitted without much regard to the negative utility of disclosure. A new awareness in both the public consciousness and the legislative agenda threatens to impede technological progress before solutions are adopted to address it.

Data is the raw ingredient in so many of today's innovations. AI is only as powerful as the data it is trained on. The more fine grained the data points and the more correlated the signals, the more predictive the models can become. Whether you are working to create new treatments for disease or the next great product, data about people is going to be vital to the analytics needed to bring it to market.

There lies the conundrum. For humans to benefit from advanced analytics and AI, it is going to require humans to become the subject matter. And to be abundantly clear, no one wants to be the data subject being tracked by unknown actors in countless databases. The future of AI needs to balance the predictive value of personal data with the privacy of individuals.

The design of a modern information sharing approach must therefore first consider the possible impact on the subjects represented by the data to be shared. Who actually owns personal data? Who controls it? Who benefits from utilizing it? What harm can come from hoarding it?

The notion of personal data rights is a rapidly evolving field with both moral and legal components. For example, within the realm of healthcare, personal health information can unlock life-saving insights that can be out-of-reach if safe sharing methods are not available. For technology to keep pace with these dynamics, new mechanisms need to be designed, developed, and implemented to establish a new social contract regarding personal data.

Connecting Protected Data

One of the major challenges of working with personal data is how to protect it. Locking it away simply precludes the discovery of any insights. Freely copying it exposes far too many data privacy and security risks to both the data subjects and the data owners. The critical need in the Protected Data Age is to link data assets at an individual level without risking exposure of the underlying identities.

There are several approaches to protecting privacy while sharing data. The techniques that attempt to control the disclosure of data are broadly referred to as Privacy Enhancing Technologies, or PETs. Even basic encryption can be considered a baseline PET.

Several recent advancements in the protection of privacy include fully homomorphic encryption and secure multi-party computation. These address more specific use cases and more nuanced protections. For example, being able to obscure underlying data values while still being able to compute aggregate functions can enable interesting analytics in a privacy preserving manner. However, the primary use-case presented here is the **linkage problem**, which has specific requirements that most PETs do not address.

The linkage problem can be defined as:

Given, two independent entities (public or private) are each managing a dataset about individuals. The understanding of an individual's identity is achieved using various identifiers such as name, postal address, email, and/or social security number. However, these components of personal data are sensitive and are tied to personal privacy rights, regulatory restrictions, and/or ethical handling concerns.

How to enable the 2 independent entities to share the understanding of the individuals in common between the 2 datasets without sharing any personal data and without inadvertently allowing reidentification of those individuals not in common (i.e., outside of the desired intersection)?

Some real-world examples of the linkage problem (where data privacy and data sharing are equally critical) include:

- Retrieving Protected Health Information (PHI) associated with a patient from another health system
- Detecting fraud or anti-money laundering activity between banking institutions
- Researching rare disease treatments by gathering longitudinal views of patient data
- Detecting duplication in voter registration databases
- Anonymizing contact tracing interactions for viral exposure monitoring

The linkage problem causes many challenges when it comes to data collaboration. Whenever records with individuals as the data subject are joined together, solutions commonly used today require that one party needs to fully trust the identity of subjects with the other party. Sometimes that is due to direct, clear text data sharing. Other times, pseudonymization techniques are used. It is important to note that pseudonymization enables reidentification when combined with additional information and is defined clearly as "personal data" under GDPR regulations in the European Union. Several categories of PETs protect

the privacy of data but still serve as a form of pseudonymized data when treating identifying data.

Partitioned Knowledge Orchestration

An emerging technique called Partitioned Knowledge Orchestration (PKO), which belongs within the family of secure multi-party computation technologies, is specifically designed as a solution for the linkage problem. It is particularly useful for data sharing operations where the disassociation and confidentiality of identity is critical.

This approach's defining feature is deliberate fragmentation of complete information that could lead to identification or re-identification. This partitioning is created during a careful orchestration of transactions performed by at least three independent actors. Of the three actors, two will be data controllers that compute identity hashes, and one will be a blind facilitator, which receives identity hashes from each data controller for comparison but never receives the one-time-use formula used to build the hashes.

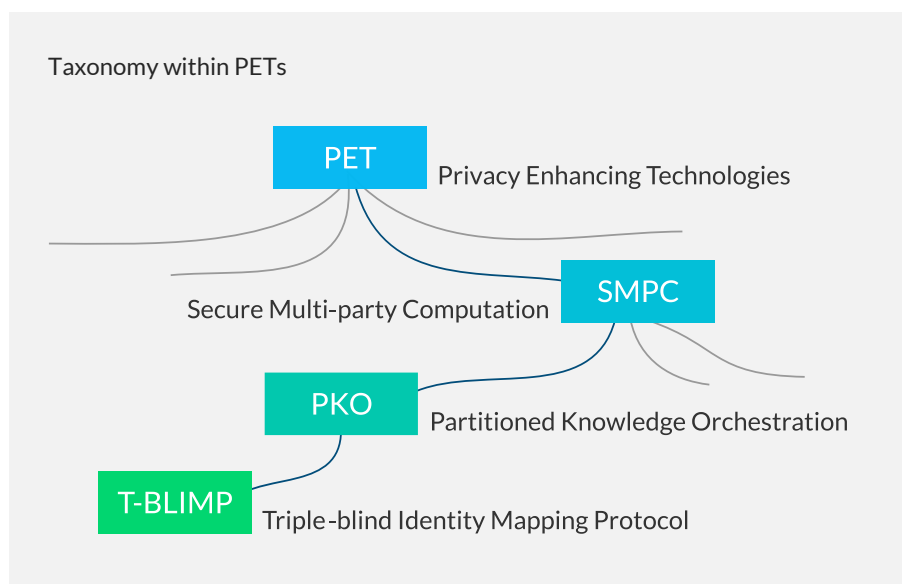
Several observable attributes of Partitioned Knowledge Orchestration can satisfy the design goals of data sharing operations that are vulnerable to the linkage problem.

- Zero-trust required of exchange partners including actors in facilitating roles, since no other party ever receives information that can lead to re-identification. This forms a stronger information security stance that cannot be compromised due to never enabling exposure. There are no data artefacts or exhaust from an exchange that can be retained to leverage for re-identification attacks.
- Zero change in data custody of identifying information before, during, or after the sharing transaction. This is vital to establishing provenance and data subject permissioning over data assets without back-door allowances for data sharing operations.
- Zero obfuscation of the collaboration methodology including cryptography, matching rules, or expected outcomes. This leads to faster adoption of the PETs and higher levels of trust among government entities, private industry, and the public at large. New cryptography poses higher risk of undiscovered vulnerabilities. A transparent approach is superior to an opaque approach when cryptanalysis and wide-spread acceptance is required.
- Zero limitations on topology of the exchange network promote an open environment that can

cross technological, industrial, and geopolitical boundaries. Since the data intersection architecture is completely distributed and decentralized, there is no limiting factor on capacity. Every transaction enrolls 3 actors: 2 encoding nodes and 1 comparison node between them. Each of these actors can be commissioned and decommissioned at will including once per transaction. This approach is a natural fit for global scale and open architecture.

A New Mechanism for Linking Protected Data

KarlsGate has developed a practical implementation of the Partitioned Knowledge Orchestration technique called Triple-blind Identity Mapping Protocol (T-BLIMP), specialized for linking identity information at scale.



The orchestration procedure of T-BLIMP follows a simple sequence.

1. The orchestration of events starts with managing single-use cryptographic keys controlled solely by the trading partners.
 - a. Each party picks a seed value (a large random number).
 - b. A shared secret is derived between the 2 partners via the selected ephemeral key agreement protocol and is never shared with anyone else.
 - c. Identifier values are extracted from the raw data, e.g., email address.
 - d. Single-use Hash Tokens are assembled from the identifier values + the seed values + shared secret + the selected cryptographic hashing

algorithm. (Hashed Tokens are different for every trade.)

2. The next step enlists a third party, the facilitator, to stand between the trading partners to perform neutral and naïve adjudication of hash collisions.
 - a. Each trading partner transmits only Hashed Tokens to the Facilitator to perform linkage.
 - b. The Facilitator compares the lists of Hashed Tokens and sends back a signal on each matching value directly to the trading partners. (The Facilitator never knows the math used to compute Hashed Tokens.)

KarlsGate has developed the Triple-blind Identity Mapping Protocol [T-BLIMP] to address the current weaknesses in identity exchange. This patented technology significantly increases the privacy protection of identifying information over current methodologies of sharing data. In this form of exchange, identities in one data set can be linked or “mapped” to the matching identities within another data set. T-BLIMP will operate on any two data sets that contain a common unique identifier. Examples would include postal address, email address, IP address, mobile advertising ID, or even more sensitive identifiers such as government identifiers.

The principle use case is to induce a shared understanding of specific sets of people without revealing the identities of those individuals. As complex as the fields of cryptography and information security can be, the implementation of this data protection scheme can be understood by analyzing the process step by step.

HASH FUNCTION

A foundational component of this methodology is to employ a cryptographic hash algorithm. A cryptographic hash algorithm is a function that exhibits several useful qualities:

- The same input always produces the same output.
- It cannot be reversed to produce the original input from the output.
- Different input values are extraordinarily unlikely to produce the same output.

Instead of comparing the values of the identifiers directly, one can compare the values of the output of a hash function. When the hash output values are the same, then it can be implied that the input values were the same. Through these stand-in yet irreversible hashes, the personally identifiable information can be protected while comparing between two data sets.

There remain a few ways that identity could still be leaked even when an adept hashing procedure is employed. The additional layers of complexity in this

mechanism are focused on preventing workarounds to the protection afforded by hashing.

To illustrate how this works, here is an example of values transformed using SHA-1, a common cryptographic hash algorithm:

Input:

tester@example.com

Output:

03055F3F1E7D343D6951533DF05E965BDF8C3F94

The modified hash output will be:

0778C8BF7390ED36B67614D5A9AA7F33FF991F46

Since a hashing function inherently produces a stable output, hash values could be stored in lieu of the original values for later comparison. If a bad actor hoards these hash values, a lookup table can be generated that will enable a third party to link identifiers to hash values later without the knowledge or consent of the original data trading partners.

RANDOM SALT

A solution to this vulnerability is to prompt each participant to generate a dynamic prefix to muddle the hashing function. This technique of introducing entropy to the hash function is referred to as ‘salt.’ Salt serves the purpose of making the resulting hash value dependent on an additional piece of information beyond the input value (i.e., the identifier). If different salt values are used for each execution, the hash values produced will be different too. Properly salted hash values cannot function as a persistent identifier since the identifier is transformed to a different value each time. A random number generator fits the requirement to produce a fresh and unpredictable seed value for the salt on demand.

This protocol requires that each time the mapping procedure is executed, participants must create a fresh random seed individually, which is then used to construct a salt value. Ultimately, this salt value is prepended to each identifier before hashing. Since the output hash values will be altered by this random influence, the resulting values have no utility for future comparisons. A participant can ensure that a truly dynamic salt value is being used by contributing their own random seed as a component of the salt value used in hashing.

By having multiple independent sources for random values that comprise the full salt prefix, it ensures that each participant can trust the salt value has not been reused from an earlier execution.

For example:

Party A contributes random seed: AAAA

Party B contributes random seed: BBBB

The salt value would be: AAAABBBB

The hash input value for tester@example.com should instead be:
AAAABBBBtester@example.com

The modified hash output will be:
0778C8BF7390ED36B67614D5A9AA7F33FF991F46

The next loophole to close is the case where one trading partner is intending to record the salt value to use with another data set. Each participant would have all the information needed to construct the salt value and could compute hashes to map against the unmatched identities from the current execution at a future time.

NEUTRAL FACILITATOR

To close this loophole, a neutral facilitator is selected to receive all the hash output values. The role of the facilitator is to act as a referee and to answer only a simple question from each trading partner: "Does this hash value exist within the other trading partner's data set?"

This keeps the trading partners from having direct access to the hash output values and only receiving a Yes or No response concerning the match of each identity. Trading partners do not receive unmatched hash values and cannot leverage them for future mapping attempts. Moreover, the facilitator will also contribute a random seed to the salt. This is done to ensure that executions are not replayed with a previously used salt value. As the facilitator is actively participating in the mapping, it has a vested interest in demanding its own source of entropy.

Another advantage of working through a neutral facilitator is to prevent inadvertent identification of a single identity. Anytime a mapping execution yields an exceptionally low occurrence of matches, it risks re-identifying an individual simply through the process of elimination or establishing a unique combination of attributes. To maintain statistical anonymity, a minimum match count is enforced to prevent this side effect. The facilitator will reply with a negative match response for every hash value whenever the number of positive matches is below the minimum threshold.

The facilitator adds several important benefits but exposes a vulnerability as well. If the facilitator were dishonest about purging all hash values after execution, it could reintroduce the same problem described earlier

by retaining old hash values along with the full set of salt components.

SHARED SECRET

An approach to prevent facilitator snooping is to keep one vital component of the salt value out of the awareness of the facilitator. If the facilitator cannot construct the exact salt value used by the trading partners, then the hash values will not be susceptible to reuse.

The next bolstering to the salt value will be something only the two trading partners can compute. A perfect method for devising a shared secret between two parties while others are observing is a public key agreement protocol. The Diffie-Hellman key exchange mechanism is a renowned example and can be utilized to enhance the salt value. To initiate the key exchange, keys are randomly generated in related pairs; one key will be publicly distributed, and one key will be privately secured.

The trading partners will leverage their respective ephemeral public/private key pairs during the setup of every execution. They each pass their public key to the facilitator, which, in turn, passes it to the other partner. The private keys should be securely stored and kept private. While in possession of a partner's public key, a shared secret can be mathematically devised. This secret value, shared only by the two trading partners, will also be appended to the salt. It is important to note that the facilitator never gains access to this component of the full salt value.

For example:

Party A contributes random seed: AAAA

Party B contributes random seed: BBBB

Facilitator contributes random number: FFFF

Party A and Party B negotiate a shared secret via public key exchange yielding: SSSS

The full salt value would be: AAAABBBBFFFFSSSS

The hash input value for tester@example.com should instead be: AAAABBBBFFFFSSStester@example.com

The modified hash output will be:
F341757CA419D6D98324E4E44A01921AE8494F08

The facilitator can serve as participant authenticator, or an established public key infrastructure (PKI) can be used. Digital certificates signed by a certificate authority (CA) can provide an additional level of verification of the participants' identities.

Automating the Orchestration

Implementing a secure data sharing protocol using T-BLIMP delivers on the data protection and safety objectives, but it is important to make the automation and coordination of a Partitioned Knowledge Orchestration accessible to all teams working with personal data. Complex cryptography can create a technical barrier to adopting such advanced practices.

The Karlsgate Identity Exchange is a data management tool designed to coordinate a vast network of connected partners and simplify the execution of PKO operations between members of the exchange.

There are 2 types of operations that the exchange can facilitate, which both leverage the T-BLIMP orchestration.

1. Remote Data Collaboration – sending and/or receiving data based on matches between two data sets without exposing identity.
2. Remote Data Integration – sending or receiving de-identified data sets without exposing identity to any third parties.

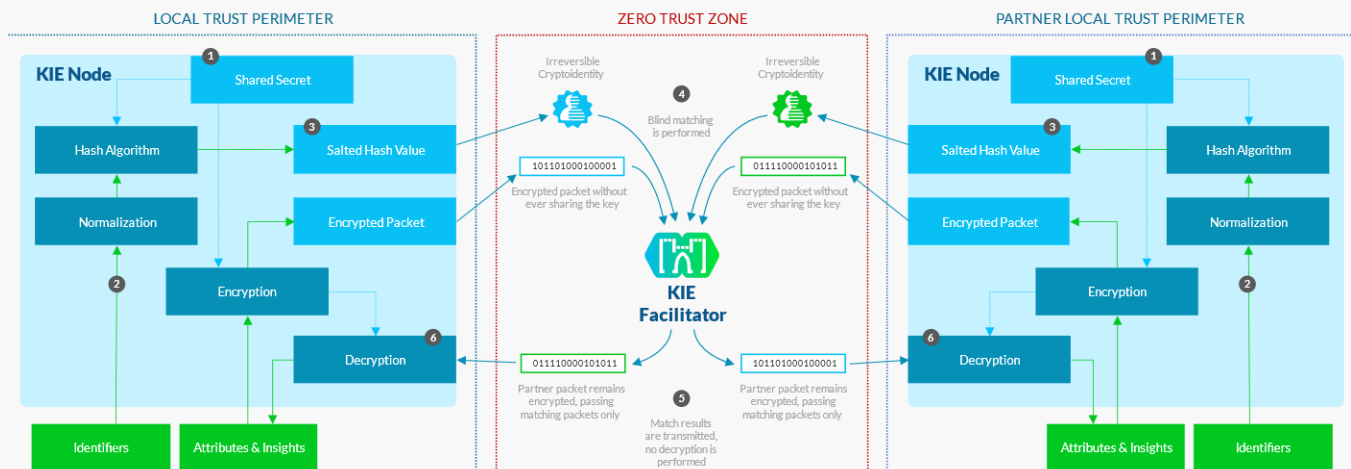
Remote Data Collaboration

Here a data transaction will be performed joining independent data sets held by 2 consenting trading partners. Abstractly speaking, the goal of a collaboration trade is to operate on the overlapping slice of a Venn diagram, while not revealing the manner in which the overlap is determined.

Both partners start with a set of identified data. Both parties must agree on the results of the collaboration which can be differing for each participant, depending on the use case. Results will be returned directly to the receiving party or parties.

This privacy-enhanced operation protects against obvious disclosure and subtle attempts to re-identify the underlying identities. None of the three participants can acquire a new identity that they did not already have direct reference to. In addition, all participants are blind to all identifying information when exchanging data. Since all the cryptographic initialization and processing is automated, leveraging the technology to solve various use cases that require joining data sets becomes trivial.

COLLABORATION WORKFLOW USING PKO



1. A transaction begins with 2 trading partners negotiating a one-time-use shared secret via a key agreement protocol. This fully automated process precedes any data processing and is discarded upon completion.
2. This shared secret serves as the salt to perform a one-way hash transformation on the normalized identifiers within the data set.
3. The resulting hash value is called a 'cryptoidentity,' which is a hash value that is secure against re-identification attempts from anyone other than the original trading partner with the same shared secret.
4. Which brings the next important step: facilitation. Since cryptoidentities are only vulnerable to the 2 trading partners, a third-party is enjoined into the transaction to act as a blind detector of hash collisions.
5. Facilitators lack the required pre-knowledge to reverse engineer cryptoidentities, so the match function can be performed without leaking identity.
6. A flow of insights protected with end-to-end encryption follows for each matching identity.

Some of the applications of a non-disclosure match capability include overlap analysis, data enhancement, longitudinal patient studies, financial fraud checks, and consent verification. The ability to perform at scale with minimal preparation adds to the versatility of the technique.

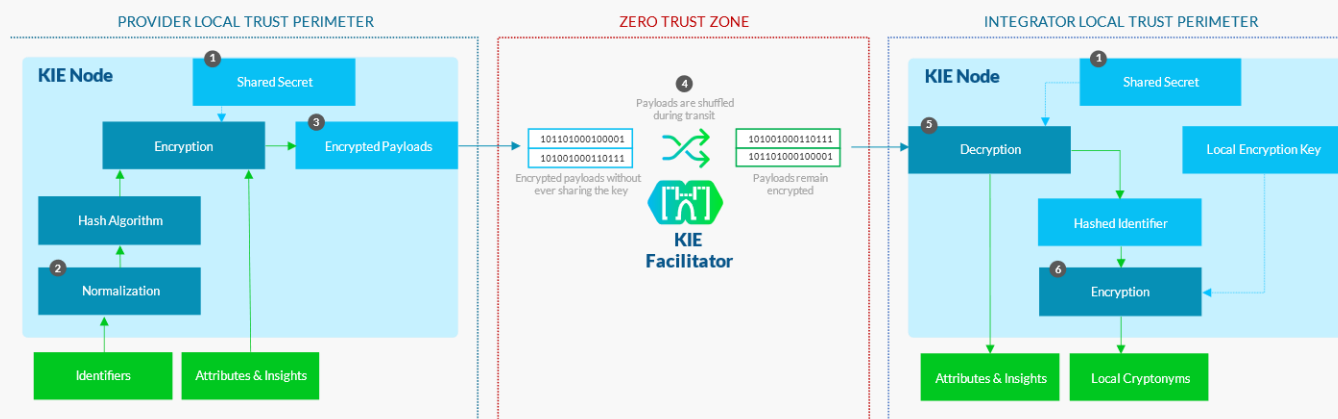
Remote Data Integration

Here a data transaction will be performed transmitting a single data set from one partner to the other. The

receiving party in an integration trade need not start with any identified data. This type of operation also leverages the privacy and safety measures afforded by T-BLIMP to ensure that the delivery mechanism cannot expose an attack vector to re-identify the traffic during the transaction.

One partner starts with a set of identified data. Both parties must agree on the results of the integration where one party will come into possession of de-identified data elements.

INTEGRATION WORKFLOW USING PKO



1. A transaction begins with two trading partners engaging in a one-time-use shared secret negotiation via a key agreement protocol. This fully automated process occurs before any data processing and is discarded upon completion.
2. The shared secret acts as the encryption key for an encryption transformation on the hashed normalized identifiers within the dataset.
3. The resulting hash values from the identifiers are combined with the attributes and then encrypted to form encrypted payloads ready for transit.
4. The subsequent critical step involves facilitation, where a third party is introduced to transfer and shuffle the encrypted payloads, adding an extra layer of protection against potential re-identification risks. Facilitators lack the necessary pre-knowledge for reverse engineering the applied encryption ensuring security throughout transit.
5. The encrypted payloads are decrypted upon reception by the partner using the same shared secret as used earlier.
6. Hashed identifiers are promptly re-encrypted using the local encryption key to generate the partner's unique, local cryptonyms, ensuring privacy compliance and mitigating re-identification risks.

Remote data integration brings privacy-enhancement and automation to data curation and consolidation efforts. Data sets can be easily de-identified, shuffled, and delivered with set-and-forget automation. Like with collaboration, no third-party can gain access to identity information.

Some of the applications of a non-disclosure integration capability include data consortiums, anonymized clinical trials, advertising cohort creation, and collecting multi-sourced AI training sets.

CRYPTONYMS

To deliver a de-identified payload that retains its ability to be matched deterministically, a specialized mechanism is needed to store pseudonymized identifiers following an integration trade. These derived values are resistant to reversal and are automatically computed during a remote data integration operation.

Integrating and consolidating data from multiple sources over time requires a unique and stable match key. Cryptonyms can serve this purpose, plus allow for

future trading activities. Not only are cryptonyms stable (i.e., the same identifier value always produces the same cryptonym value for a given data controller), but the stored value is unique to the participant (i.e., the cryptonym value for the same identifier value is completely different at each data controller).

To illustrate:

Original identifier	Resulting cryptonym value at company A	Resulting cryptonym value at company B
jane@future-mail.com	ffYuzR93tlC5bcBq...	PkwX69p8F1FLDzKF...
jane@future-mail.com	ffYuzR93tlC5bcBq...	PkwX69p8F1FLDzKF...
beni@cloud-shock.io	tNEAyfcPs6J0nKsV...	4VHJFVbtBvciFuB5...

To keep total control of locally stored cryptonyms in the hands of each data controller, a local encryption key is stored and managed by each participant, creating their own, private key space. Local encryption keys can be stored in a file or locked in a key vault. When cryptonyms need to be stored, the local key is used to encrypt incoming hashed values into a unique value for storage. When cryptonyms are used for trading, the local key is used to decrypt the underlying hashed values to transmit in a data collaboration or integration transaction.

Unlike every other aspect of PKO operations, a cryptonym is a realized exchange of identifying information. This means pseudonymous values that can be used for re-identification purposes are shared as an output of a trade. T-BLIMP ensures secure end-to-end data transfer, so no intermediary parties gain access to pseudonymous information. However, it is important to balance the risk of transacting over information that can be leverage for re-identification.

Cryptonyms are encoded using a cryptographic recipe called a ‘scheme’. The current scheme supported by the Karlsgate Identity Exchange is the **encryptid-2023** specification with extensibility to add additional schemes in the future. Encryptid-2023 utilizes a SHA-256 hash followed by an AES-SIV cipher and a base64url encoding format.

Cryptographic Adaptability

A key aspect of working with PETs is the need for peer-reviewed cryptoanalysis to vet the safety, fidelity, and reliability of new cryptographic technologies. Since Partitioned Knowledge Orchestration reuses proven cryptography in an orchestrated manner, it has a key

advantage in institutional settings. There is no need to review, analyze, and vet new and emerging cryptographic algorithms, each with its own weakness profile. Instead, the sequenced orchestration can rely completely on proven, FIPS-compliant cryptography that can be substituted with other approved algorithms over time in a very natural evolution. Again, this is not a ‘black box’ technique but an orchestrated framework of interactions with defined actions including secret key derivation, one-way hashing, and symmetric encryption transformations.

Some common choices for the cryptographic algorithms are: Elliptic Curve Diffie-Hellman (secp384r1) for the secret key derivation covered by FIPS PUB 186-4, HMAC-SHA-384 for the one-way hashing covered by FIPS PUB 180-4, and AES-256-CBC for the symmetric encryption covered by FIPS PUB 140-2. The use of well-studied and FIPS-

compliant cryptography can accelerate adoption and application of PETs in government settings.

Flexibility and adaptability are built into the technology, which takes on special importance with the presence of quantum computing. Adopting the latest NIST post-quantum cryptography standards, such as module-lattice key encapsulation (FIPS 203), empowers T-BLIMP to keep ahead of emerging threats.

Some of the additional algorithms that are available today include, for key agreement: X25519, X448, ML-KEM, SecP256r1MLKEM768; for hashing: HMAC-SHA-3, KMAC-256; for encryption: AES-256-GCM, ChaCha20Poly1305.

With Protection Comes Empowerment

The advent of the information age has brought forth unprecedented challenges in privacy and data security. As data becomes the cornerstone of innovation, particularly in AI, the need to balance the predictive value of personal data with individual privacy has become paramount. By overcoming the linkage problem, technology can enable a shared understanding between datasets without compromising personal data.

Emerging techniques such as Partitioned Knowledge Orchestration (PKO) and the Triple-blind Identity Mapping Protocol (T-BLIMP) developed by Karlsgate offer promising solutions to this problem. These techniques allow for the linking of data assets at an individual level without risking exposure of underlying identities, thereby addressing the critical need in the Protected Data Age.

These advancements in Privacy Enhancing Technologies (PETs) are paving the way for a new social

contract regarding personal data. They enable a zero-trust environment, maintain data custody, offer transparency in collaboration methodology, and allow for an open, scalable exchange network. Data practitioners now have a practical approach to privacy-aware data connectivity that fosters operational excellence with an “adopt once, use everywhere” tool.

In conclusion, while the challenges of privacy and data security in the information age are significant, the development and implementation of innovative techniques like PKO and T-BLIMP are making strides towards resolving these issues. These solutions hold the potential to unlock the full power of data-driven innovation while preserving the privacy rights of individuals, thus heralding a new era in the realm of data sharing and privacy.

Comparisons with other PETs

The following table details the comparison of common PETs for use with personal identity linkage applications:

Encryption Using a secret key, information is scrambled until the key is re-applied	<ul style="list-style-type: none"> Only protects data in transit The recipient will decrypt back to fully identified data upon processing Full custody change
Hashing / Tokenization One-way scrambling of data that is exceedingly difficult to reverse	<ul style="list-style-type: none"> Full change of custody leads to uncontrolled, future re-identification attempts against an identity graph Data owner has no agency over copies of pseudonymous data and that is why the practice is not GDPR-compliant
Data Clean Room Full dataset is sent to a controlled and isolated environment that is typically managed by a third party	<ul style="list-style-type: none"> Each usage represents a full custody change event if cleartext identifiers are ever transmitted, which may be mitigated by Confidential Computing environments with the downside of additional complexity and attestation requirements Both parties must trust the same operator with full vetting, legally binding responsibilities, consent, and security obligations Difficult to have a single clean room service for all partners and there is no clean room-to-clean room exchange mechanism
Fully Homomorphic Encryption An advanced form of cryptography that allows analysis of data without decrypting the payload	<ul style="list-style-type: none"> Does not solve the linkage problem; while the exchanged attributes are protected, the payload can still be attacked for re-identification through joining to a known identity graph Performance problems make this technique max-out at ~5 million records practically speaking
Federated Learning Building an aggregated model constructed from multiple, localized machine learning processes	<ul style="list-style-type: none"> Does not solve the linkage problem; while powerful in building audiences and propensities in a privacy sensitive manner, it does not produce a deterministic intersection between two identified data sets, therefore limiting its use cases
Differential Privacy A process of adding randomly generated adjustments to data without introducing significant skew	<ul style="list-style-type: none"> Does not solve the linkage problem; while it can be helpful for obfuscating descriptive values, it has no purpose for directly linking one identity to another May degrade accuracy for certain use cases
Secure Multi-Party Computation - Partitioned Knowledge Orchestration A coordination to share common identities while simultaneously blocking unwanted reidentification as a consequence of interacting	<ul style="list-style-type: none"> No identifying information, including ciphertext or hashes, ever flow directly from partner to partner The facilitator is deprived of any cryptographic parameters, blinding it to any ability beyond simple comparisons No trust is required to safely connect data sets and extract insights, since no identity information is put at risk Mapping rules and cryptography is obvious to all parties

Cited Technological Concepts

Cryptographic Hash Function	https://en.wikipedia.org/wiki/Cryptographic_hash_function
Differential Privacy	https://en.wikipedia.org/wiki/Differential_privacy
Diffie-Hellman Key Exchange	https://en.wikipedia.org/wiki/Diffie-Hellman_key_exchange
GDPR	https://en.wikipedia.org/wiki/General_Data_Protection_Regulation
Federated Learning	https://en.wikipedia.org/wiki/Federated_learning
HMAC	https://en.wikipedia.org/wiki/HMAC
Homomorphic Encryption	https://en.wikipedia.org/wiki/Homomorphic_encryption
ML-KEM	https://en.wikipedia.org/wiki/Kyber
Post-quantum Cryptography	https://en.wikipedia.org/wiki/Post-quantum_cryptography
Privacy by Design	https://en.wikipedia.org/wiki/Privacy_by_design
Public Key Infrastructure	https://en.wikipedia.org/wiki/Public_key_infrastructure
Random Seed	https://en.wikipedia.org/wiki/Random_seed
Salt	https://en.wikipedia.org/wiki/Salt_(cryptography)
Secure Multi-party Computation	https://en.wikipedia.org/wiki/Secure_multi-party_computation
SHA-1	https://en.wikipedia.org/wiki/SHA-1
SHA-2	https://en.wikipedia.org/wiki/SHA-2
SHA-3	https://en.wikipedia.org/wiki/SHA-3