

Zeph: Cryptographic Enforcement of End-to-End Data Privacy



Lukas
Burkhalter*



Nicolas
Küchler*



Alexander
Viand



Hossein
Shafagh



Anwar
Hithnawi

~ **2.5** million terabytes

estimate of data generated
per day in 2020

... much of this data is **personal**





data-driven world ...

... data breaches, data misuse



You got this ad because
you're a newlywed pilates
instructor and you're
cartoon crazy.

This ad used your location
to see you're in La Jolla.

You're into parenting blogs
and thinking about LGBTQ
adoption.



... privacy laws,
user awareness

Data Privacy Landscape

Compliance

"notice and consent"

Privacy-Enhancing Technologies

"ad-hoc solutions"

Zeph: Cryptographic Enforcement of End-to-End Data Privacy




User-centric Model
for Privacy




Cryptographically
Enforces Privacy

One of Many Scenarios

“Raw Location Data”



Privacy Transformation



“Daily popular running tracks”

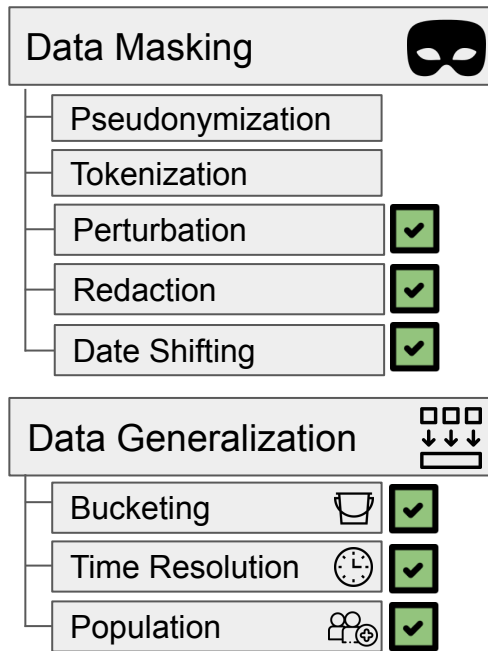


Day 6

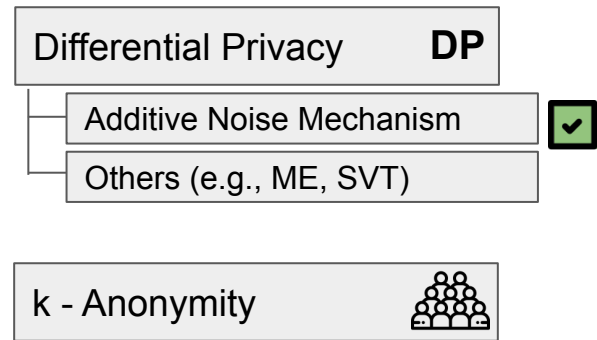
Privacy Transformations

 Supported in **Zeph**

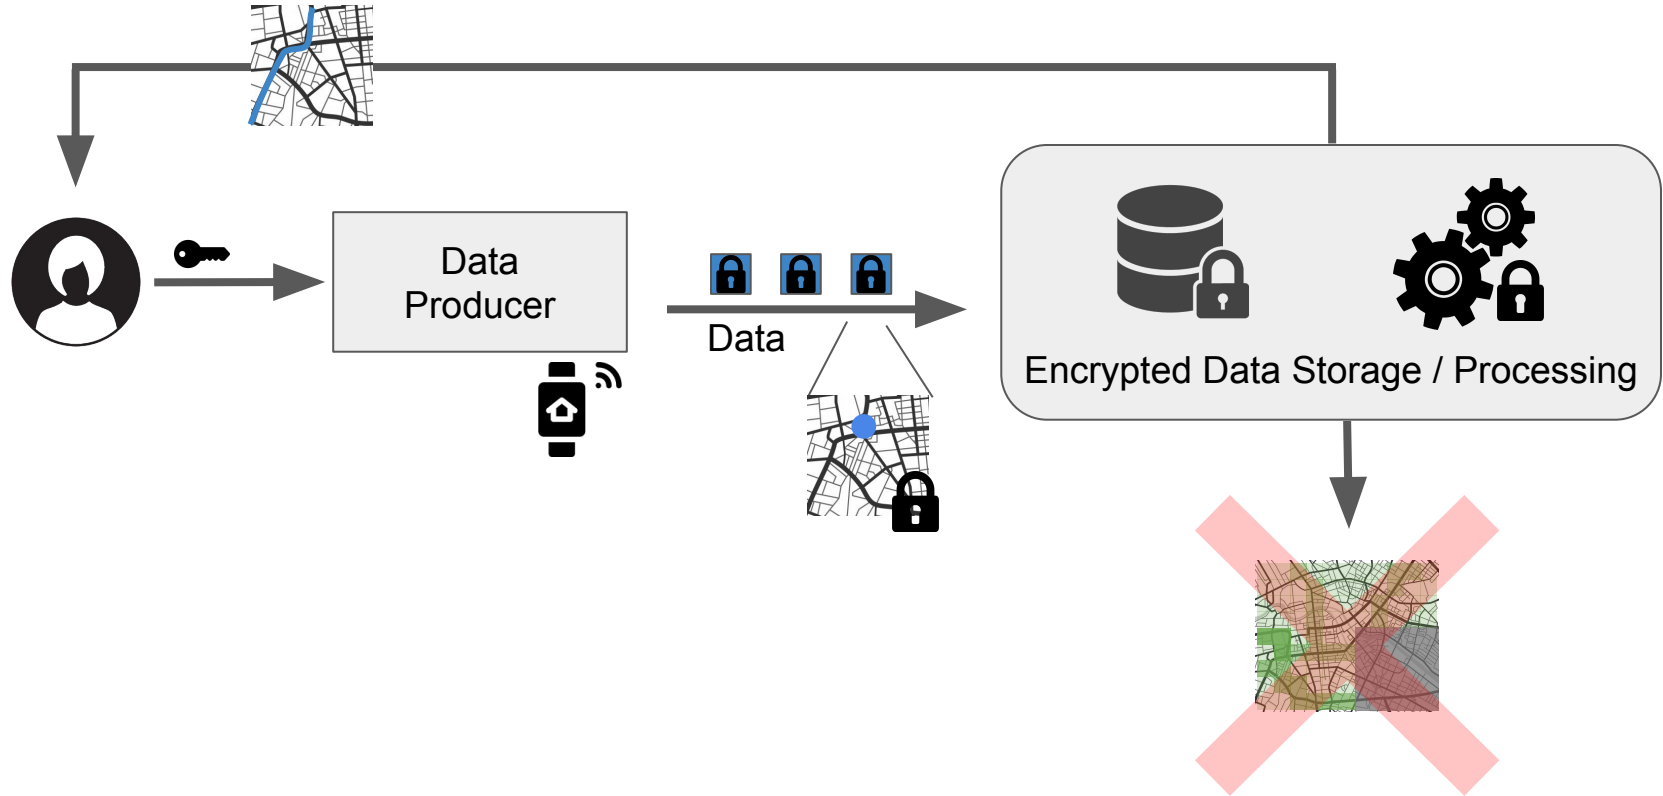
“Practical” Privacy Tools



Formal Privacy Models




Existing End-to-End Encrypted Streaming Pipeline



1. Compatibility with Existing Systems



2. Data with Heterogeneous Privacy Policies



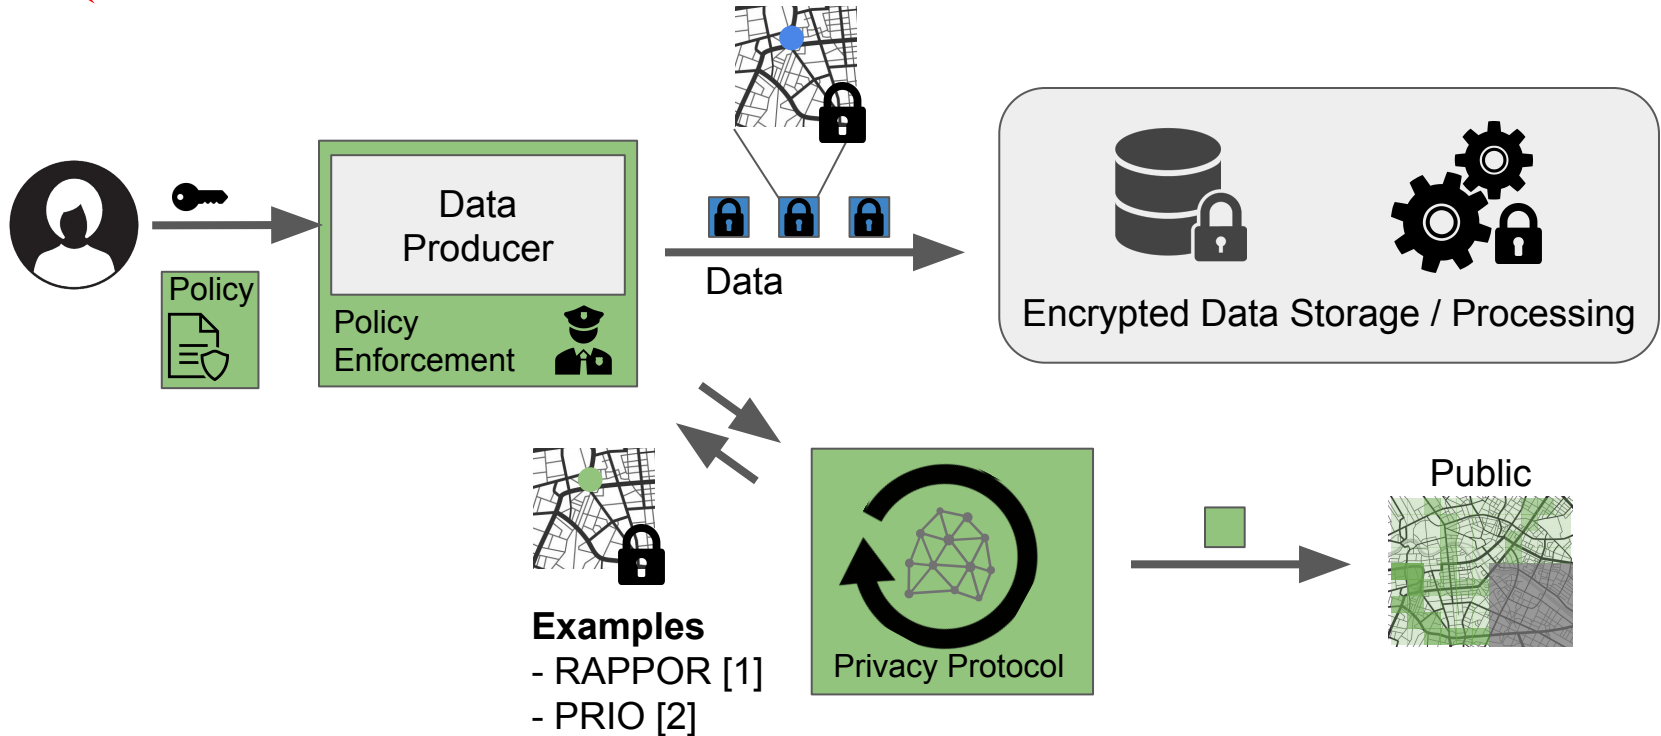
3. Allow Transformations on Encrypted Data



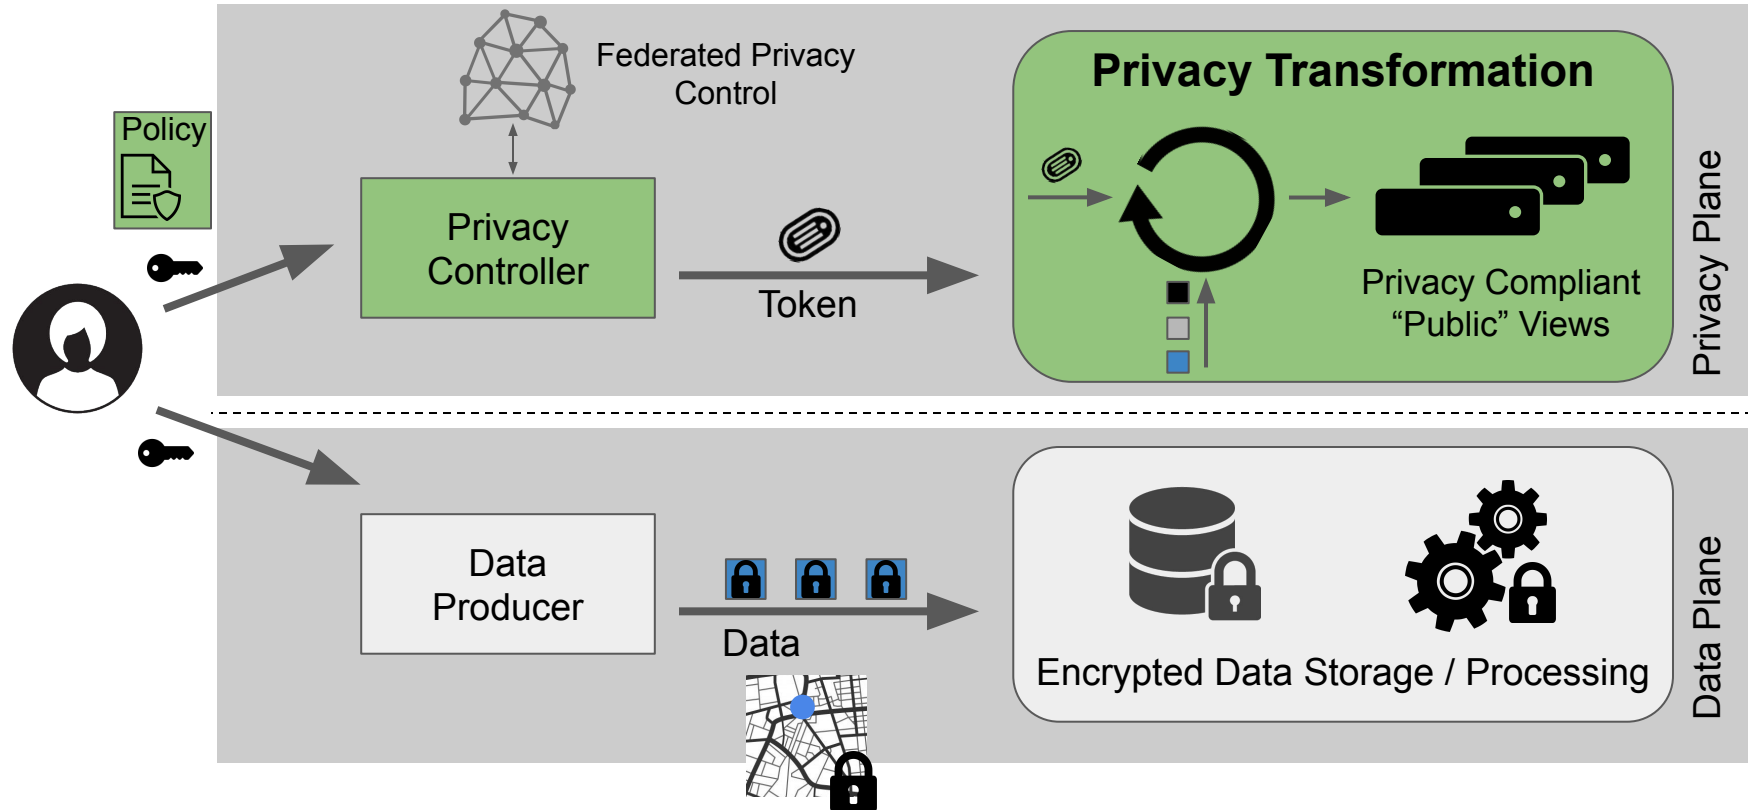
Integrate Privacy Controls into Existing Pipelines



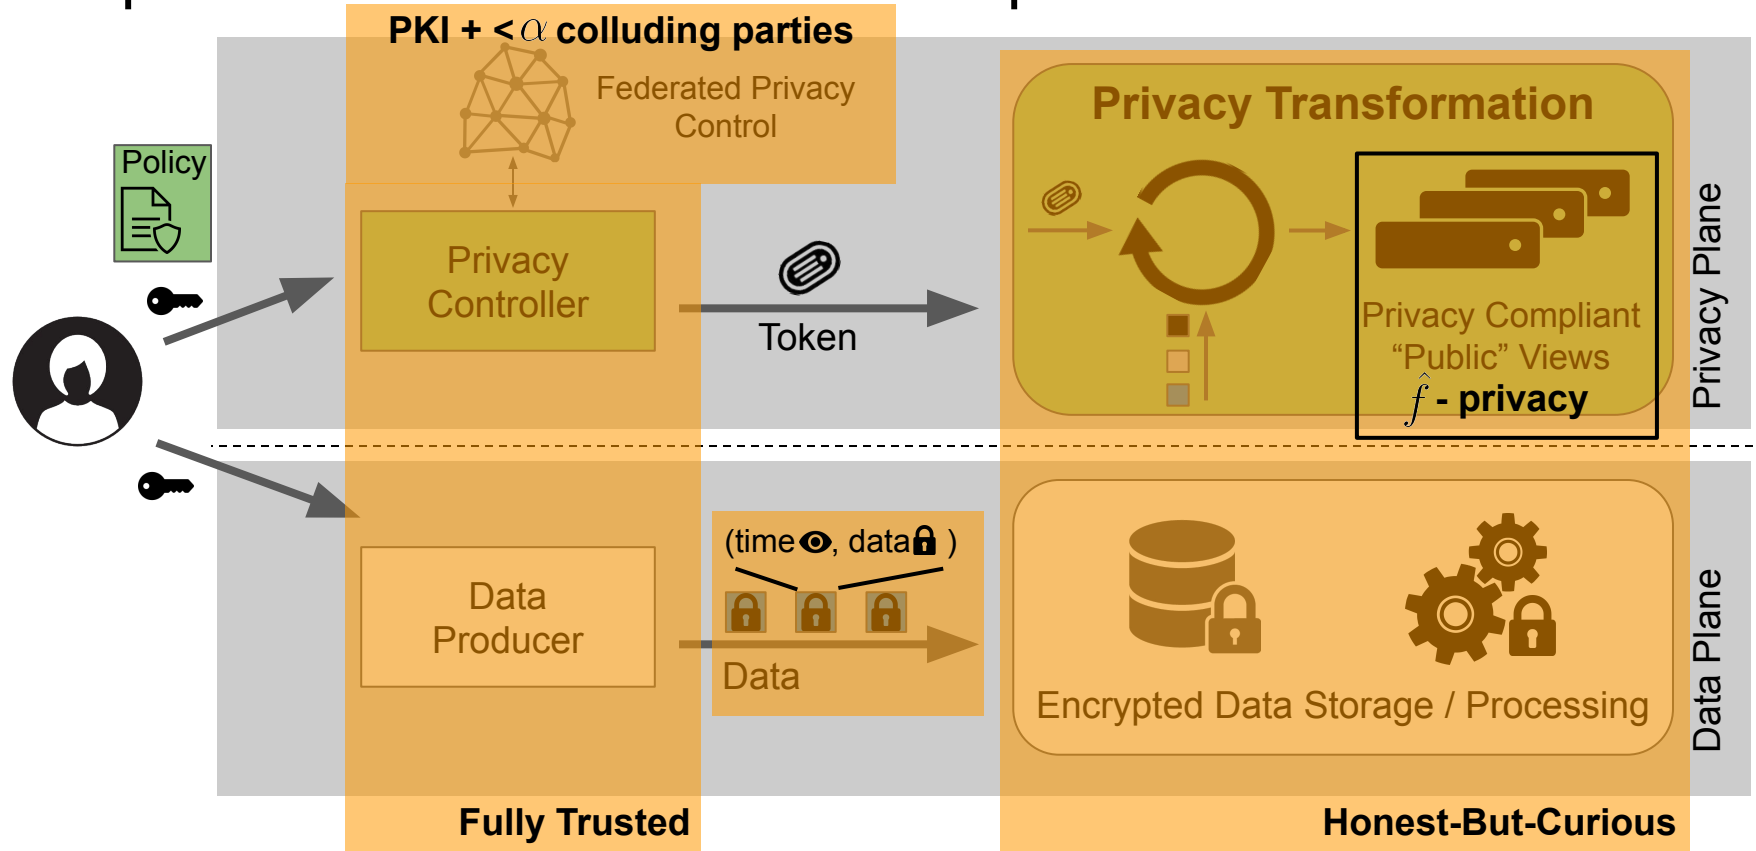
Complex and ad hoc solutions



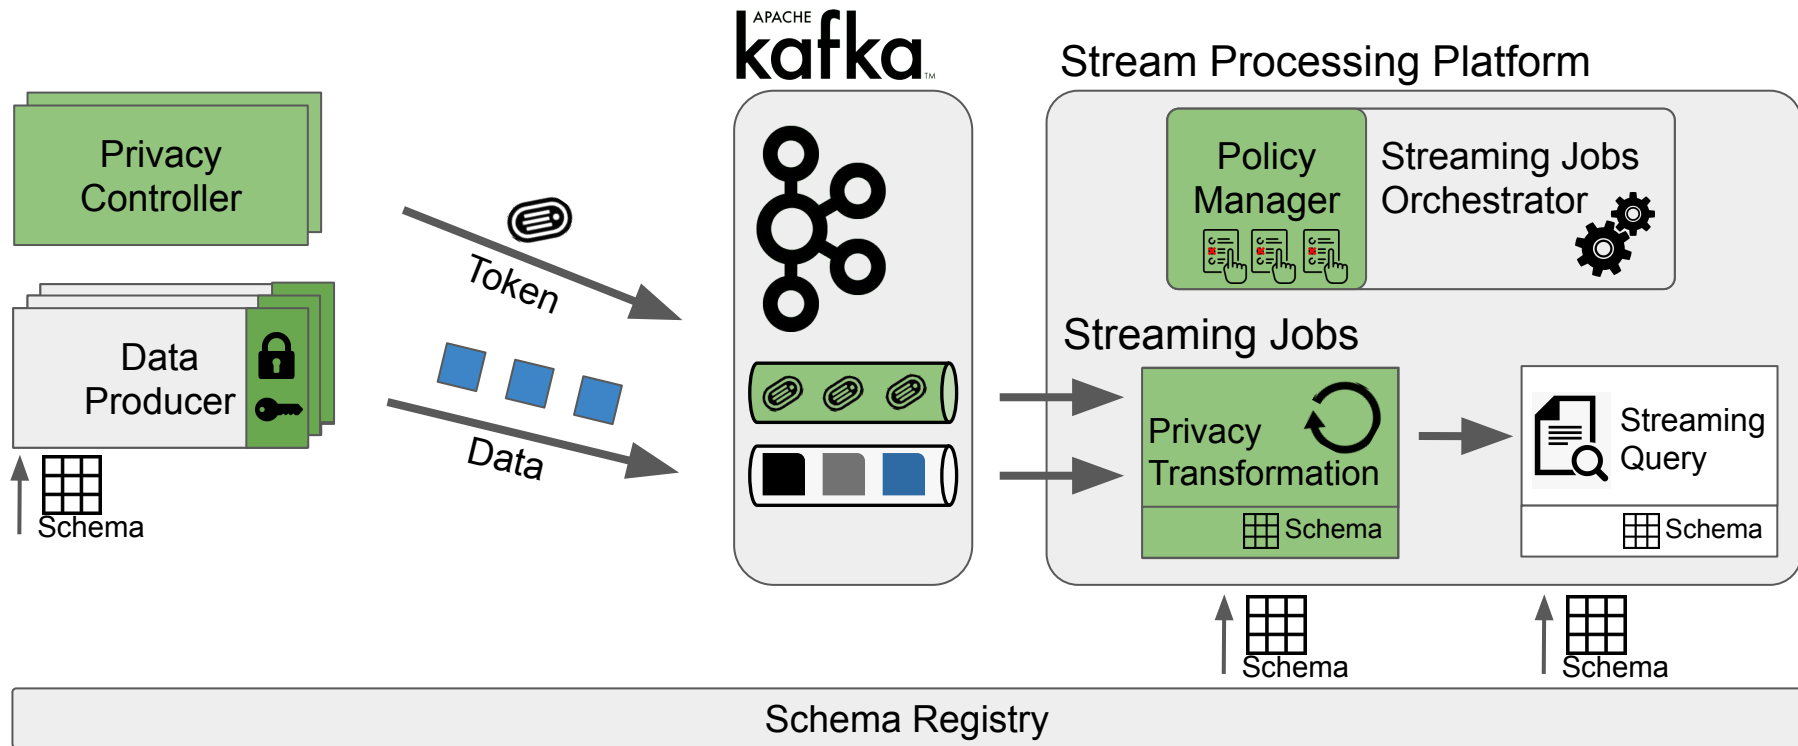
Zeph's End-to-End Approach to Privacy



Zeph's Threat Model and Assumptions



How Zeph augments existing System Designs



From Privacy Policies to **End-to-End Privacy** Challenges

Challenge #1: Keep End-User Control Simple

Privacy Preferences

- ☐ do not share data
- ☐ share data without restrictions
- ☒ share generalized views of aggregated data

End-User



Expert

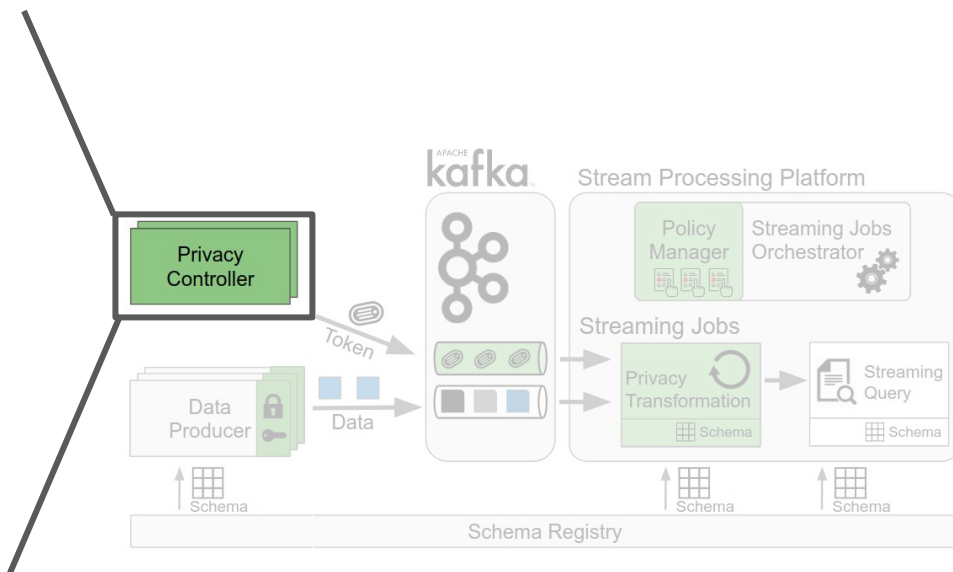
Mapping



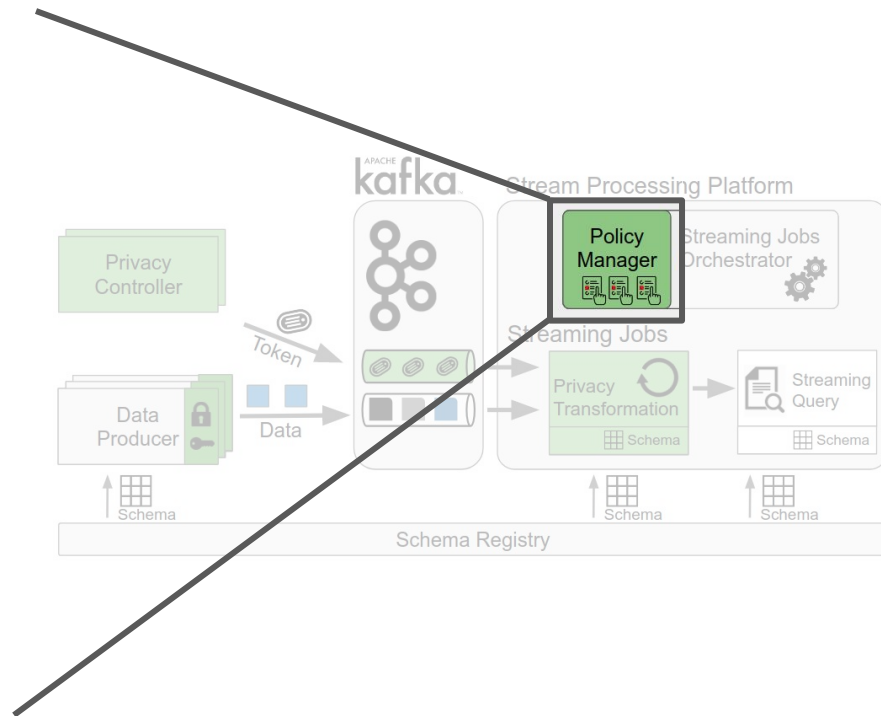

Policy Choice

Stream Schema

	Fields, Data Types
	Policy Options



Challenge #2: Organize Privacy Transformations



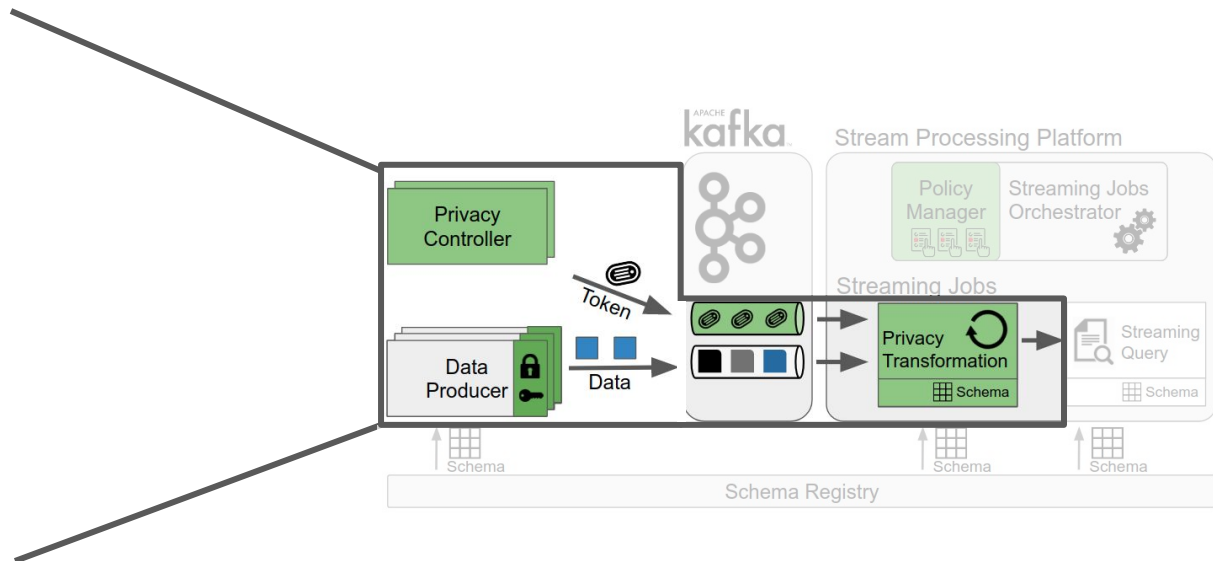
Challenge #3: Meeting Privacy Transformation Requirements

1) **Confidentiality** of data

2) Transformation **Authorization** by Privacy Controller

3) **Compute** transformation on confidential data

4) Privacy Controller is **efficient** and **independent** of data



Additive Homomorphic Secret Sharing

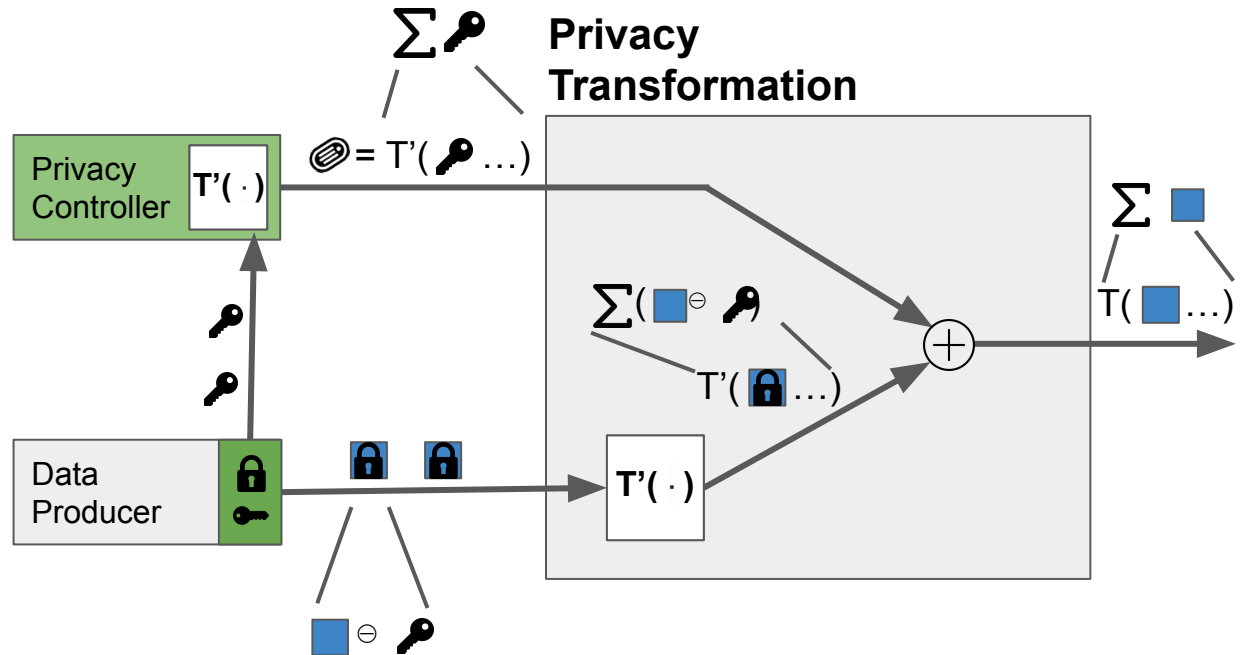
Challenge #3: Meeting Privacy Transformation Requirements

1) **Confidentiality** of data ☒

2) Transformation **Authorization** by Privacy Controller ☒

3) **Compute** transformation on confidential data

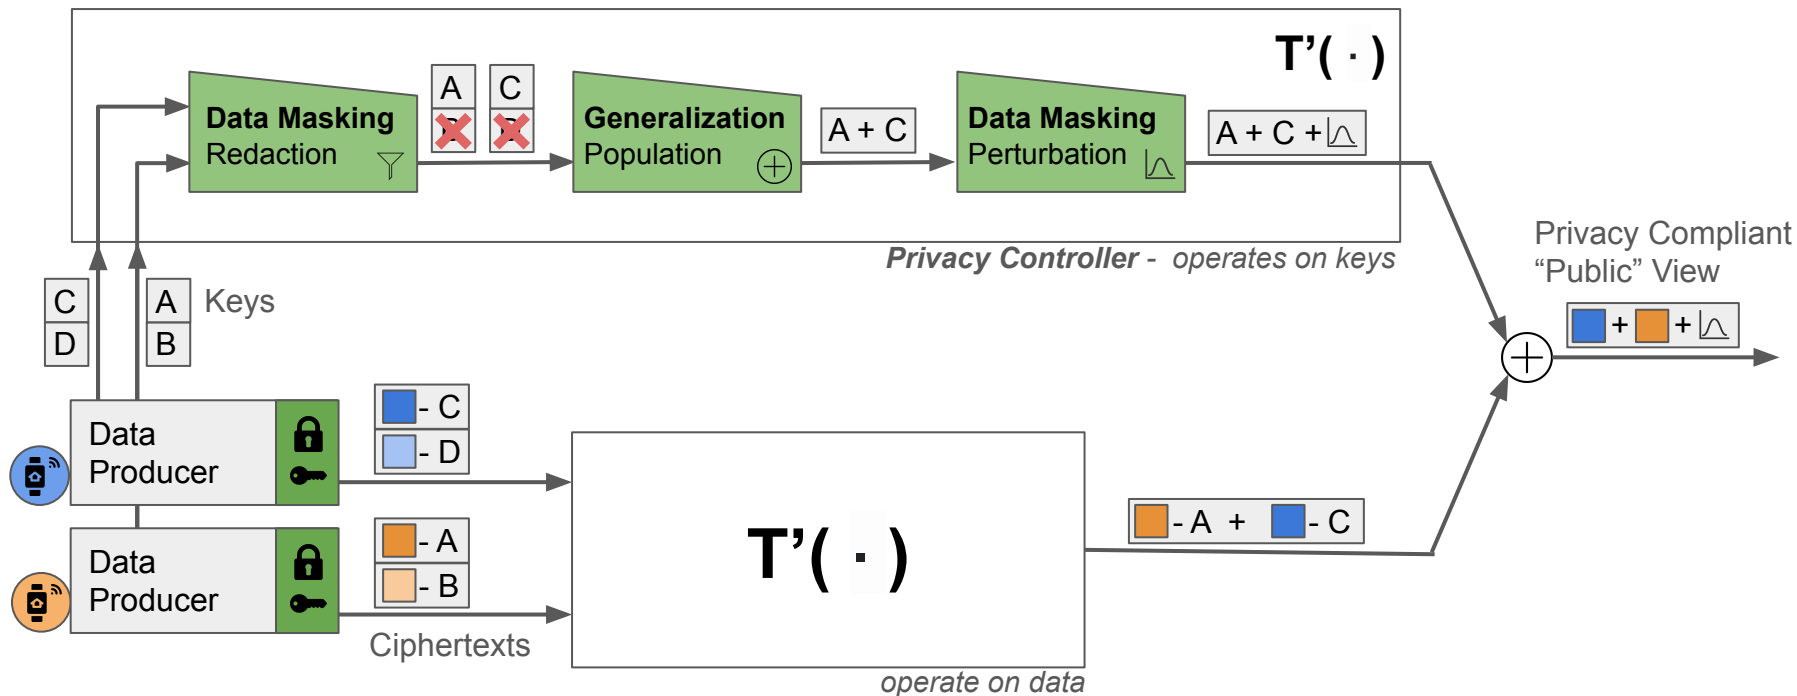
4) Privacy Controller is **efficient** and **independent** of data



Additive Homomorphic Privacy Transformations

Challenge #3: Meeting Privacy Transformation Requirements

$$T'(\cdot) = \Sigma$$



Additive Homomorphic Secret Sharing

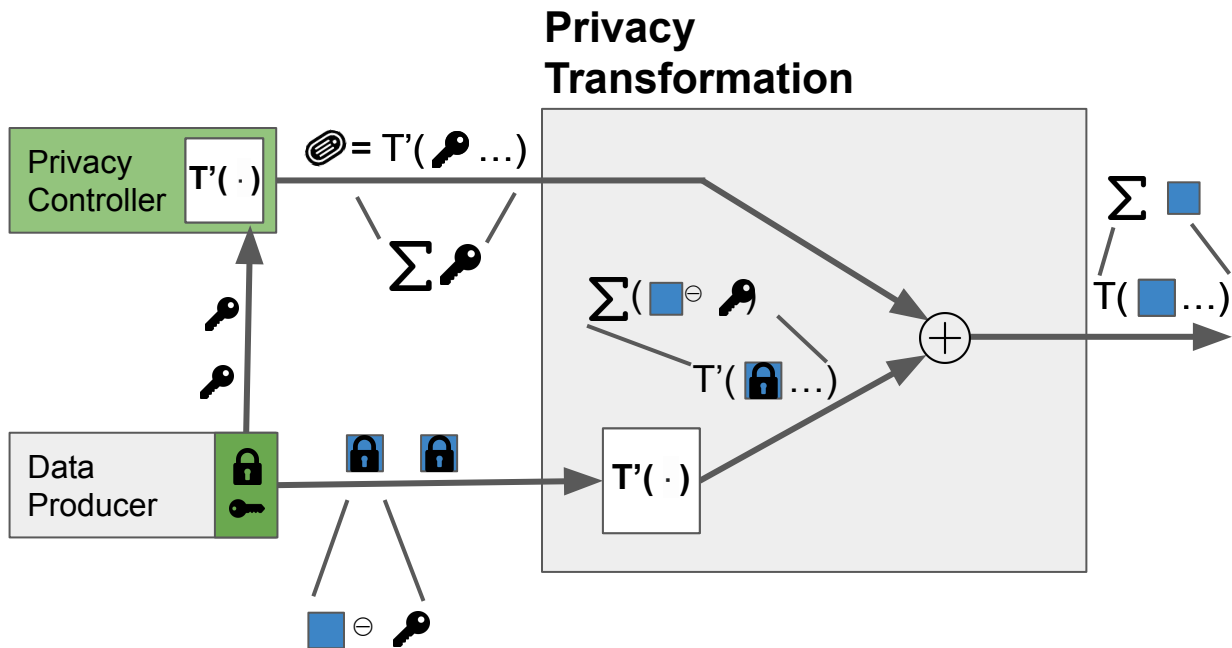
Challenge #3: Meeting Privacy Transformation Requirements

1) **Confidentiality** of data ☒

2) Transformation **Authorization** by Privacy Controller ☒

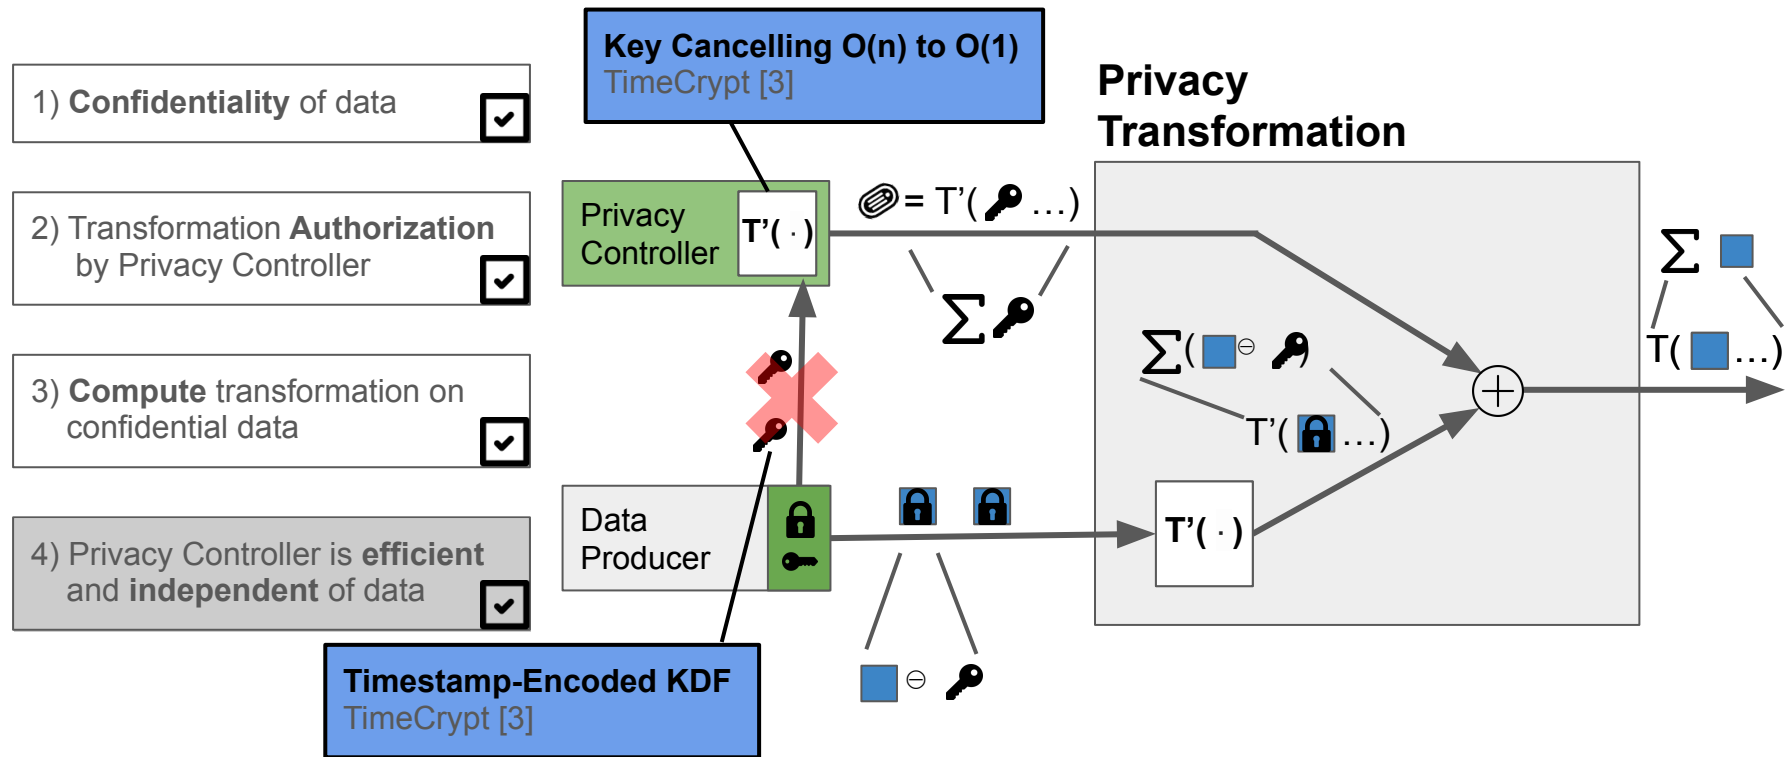
3) **Compute** transformation on confidential data ☒

4) Privacy Controller is **efficient** and **independent** of data



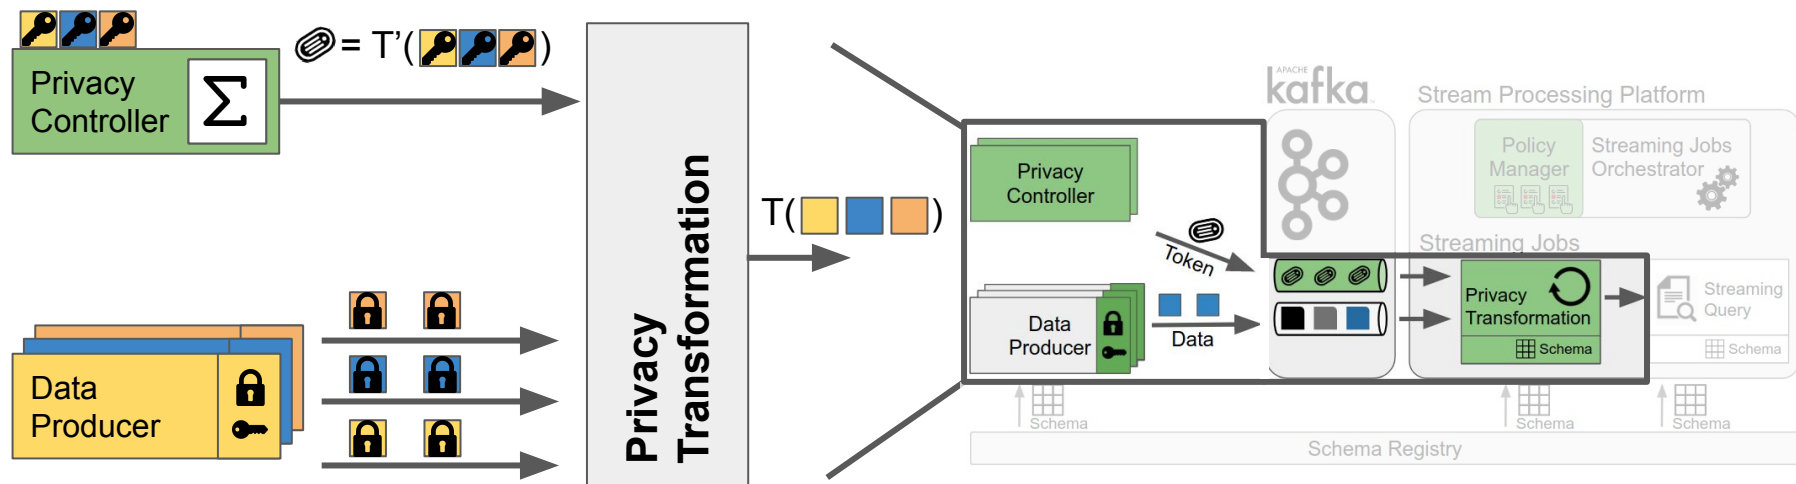
Independent and Efficient Privacy Controller

Challenge #3: Meeting Privacy Transformation Requirements



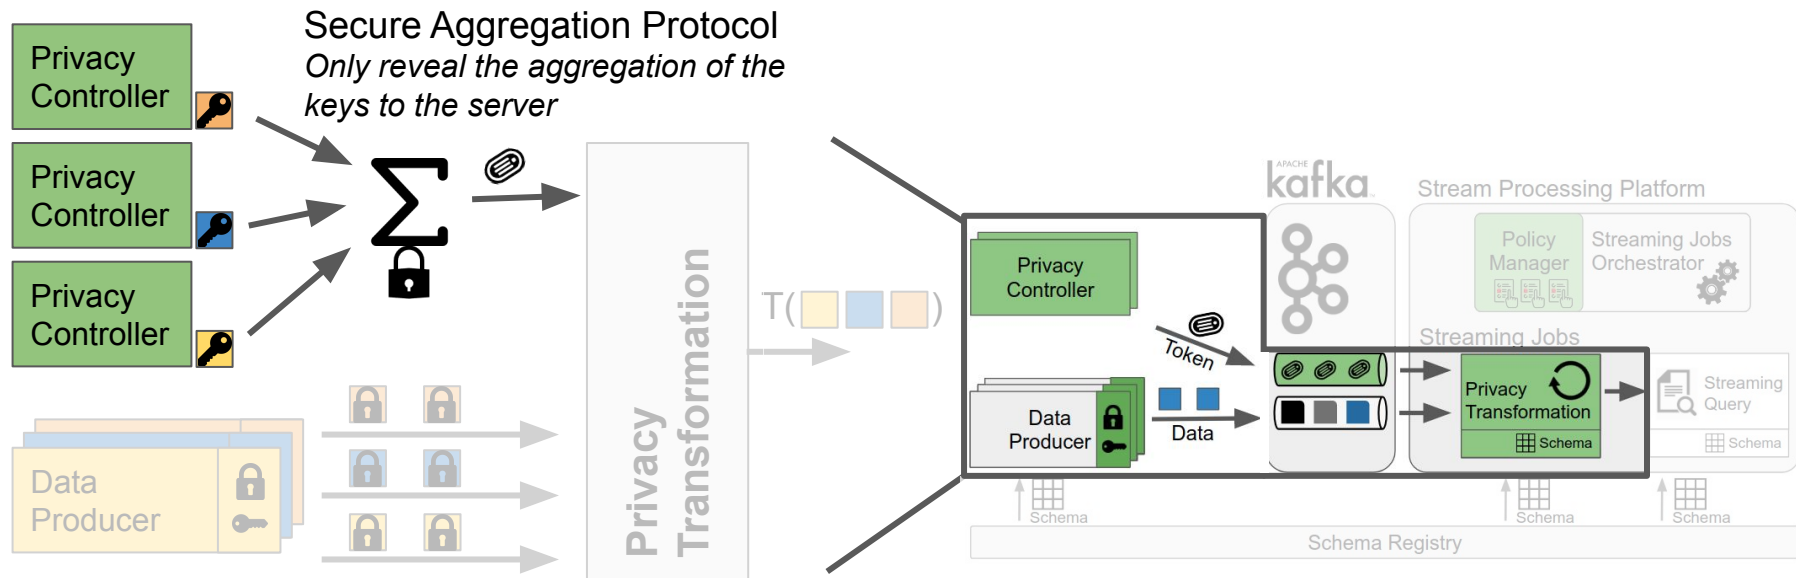
Challenge #4: Enable Federated Privacy Control

“multiple Data Producers - **one** Privacy Controller”



Challenge #4: Enable Federated Privacy Control



“multiple Data Producers - **multiple** Privacy Controllers”



Zeph Implementation and Evaluation



Data Producer & Privacy Controller



The Rust Programming Language



Privacy Transformation



Fitness App

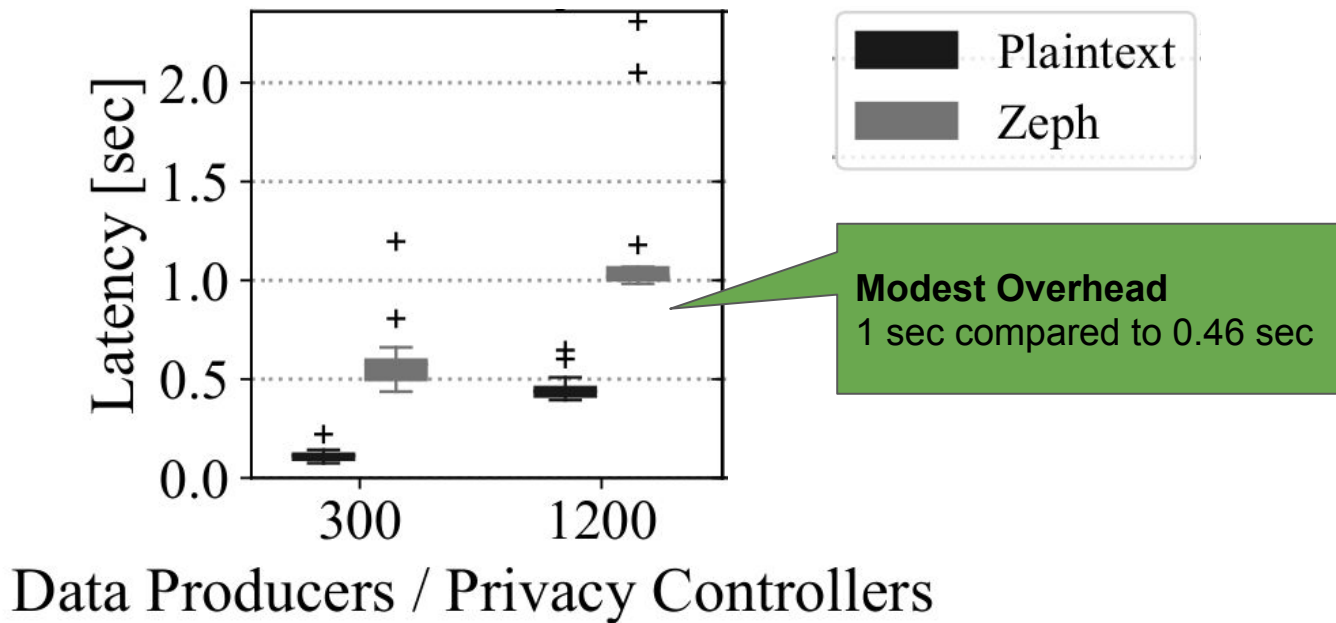


Website Analytics




Smart Car

Web Analytics: End-to-End Benchmark





pps-lab.com/research/e2e-privacy



github.com/pps-lab/zeph-artifact

Citations

- [1] Úlfar Erlingsson, Vasyi Pihur, and Aleksandra Korolova. 2014. RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. In Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS '14). Association for Computing Machinery, New York, NY, USA, 1054–1067.
- [2] Henry Corrigan-Gibbs and Dan Boneh. 2017. Prio: private, robust, and scalable computation of aggregate statistics. In Proceedings of the 14th USENIX Conference on Networked Systems Design and Implementation (NSDI'17). USENIX Association, USA, 259–282.
- [3] Lukas Burkhalter and Anwar Hithnawi and Alexander Viand and Hossein Shafagh and Sylvia Ratnasamy. 2020. TimeCrypt: Encrypted Data Stream Processing at Scale with Cryptographic Access Control. 17th {USENIX} Symposium on Networked Systems Design and Implementation (NSDI'20). USENIX Association, Santa Clara, 835-850.