Building the next generation of digital advertising with MPC

Private Computation, RWC 2022

James Reyes
Tech Lead, Private Computation

Meta

Agenda

01 Background

02 Product Market Fit

03 MPC at Scale

04 Closing & Resources

01 Background

December 2019

We believe privacy-enhancing technologies will support the next generation of digital advertising. What will our next generation ads stack look like?

Private Computation

Cryptography Driven. Started as a partnership between cryptography researchers and engineers.

Secure Multi-Party Computation. Started with using MPC w/ a semi-honest threat model.

Measurement Focused. Measurement is foundational to digital advertising.

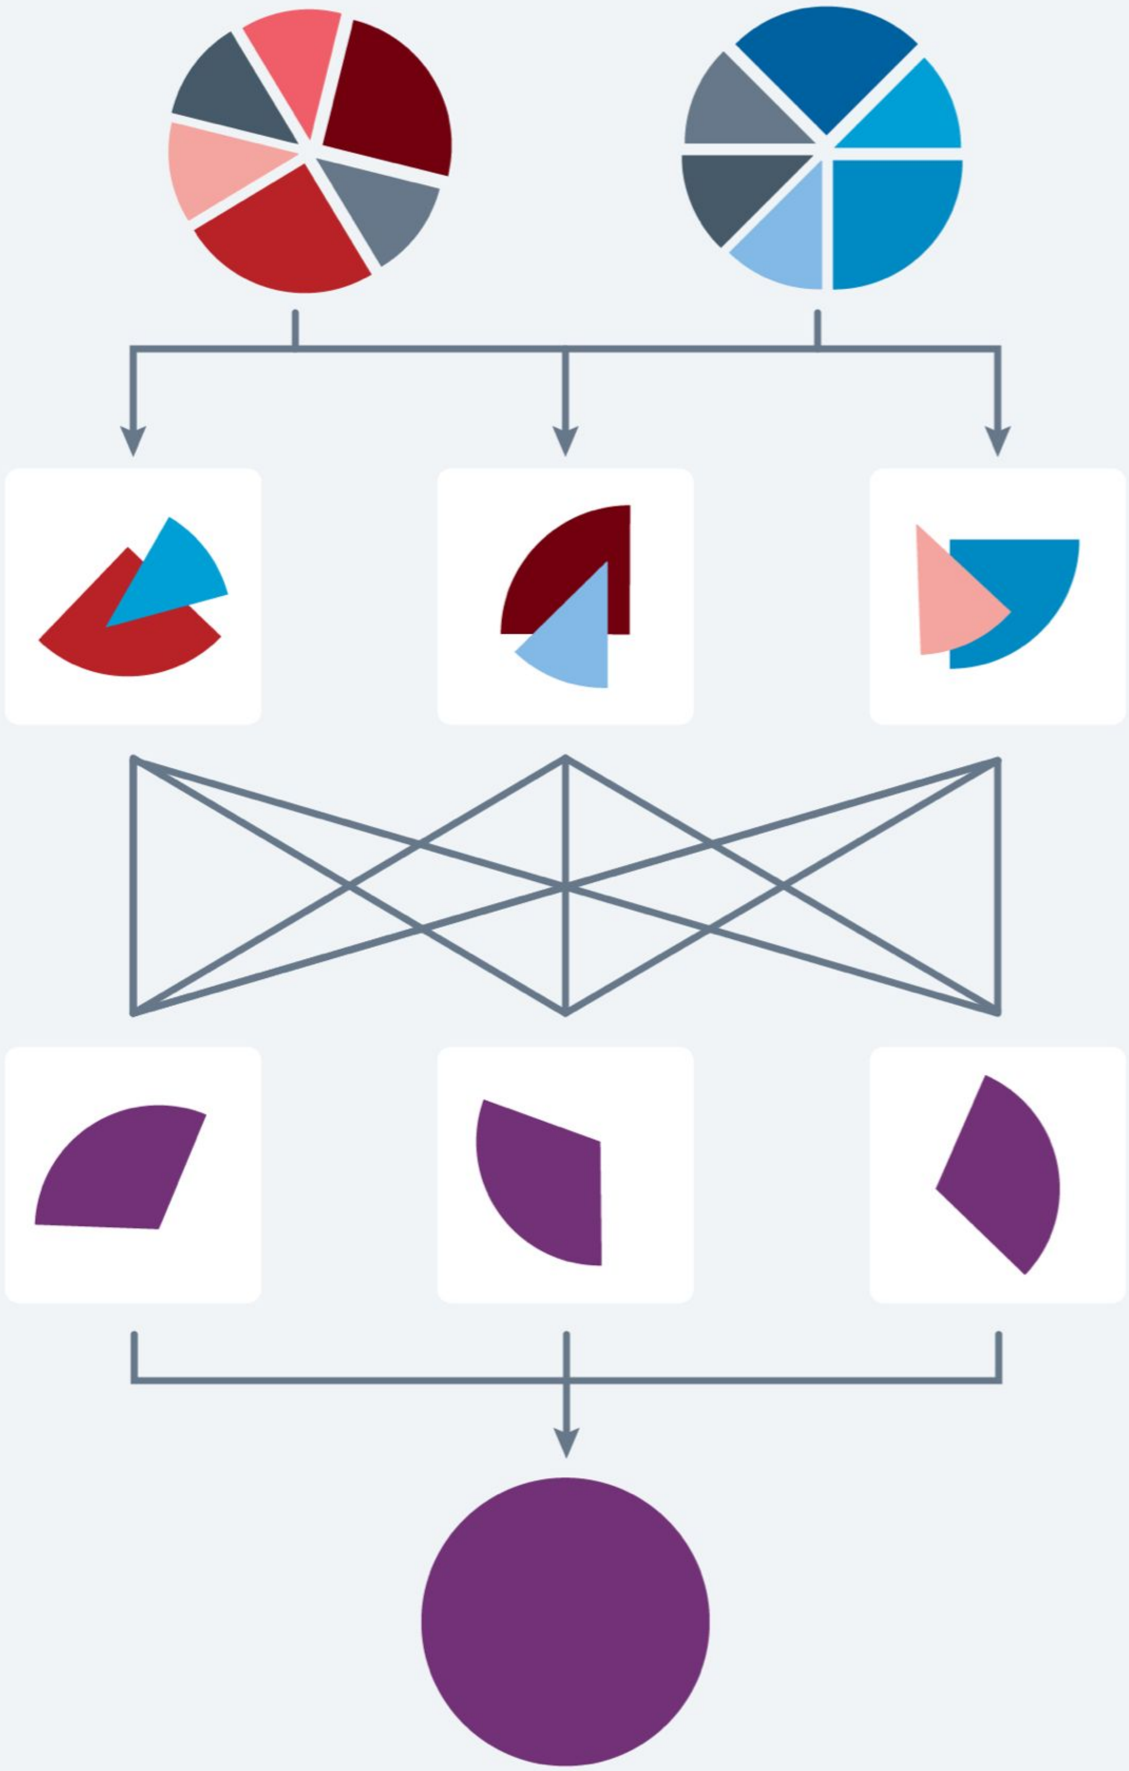

Today

- Meta has embarked on a multi-year effort to build a portfolio of Privacy Enhancing Technologies
- Cross-discipline team of over 70 engineers and cryptographers
- Two products in Beta: Private Lift and Private Attribution
- Everything is open-sourced

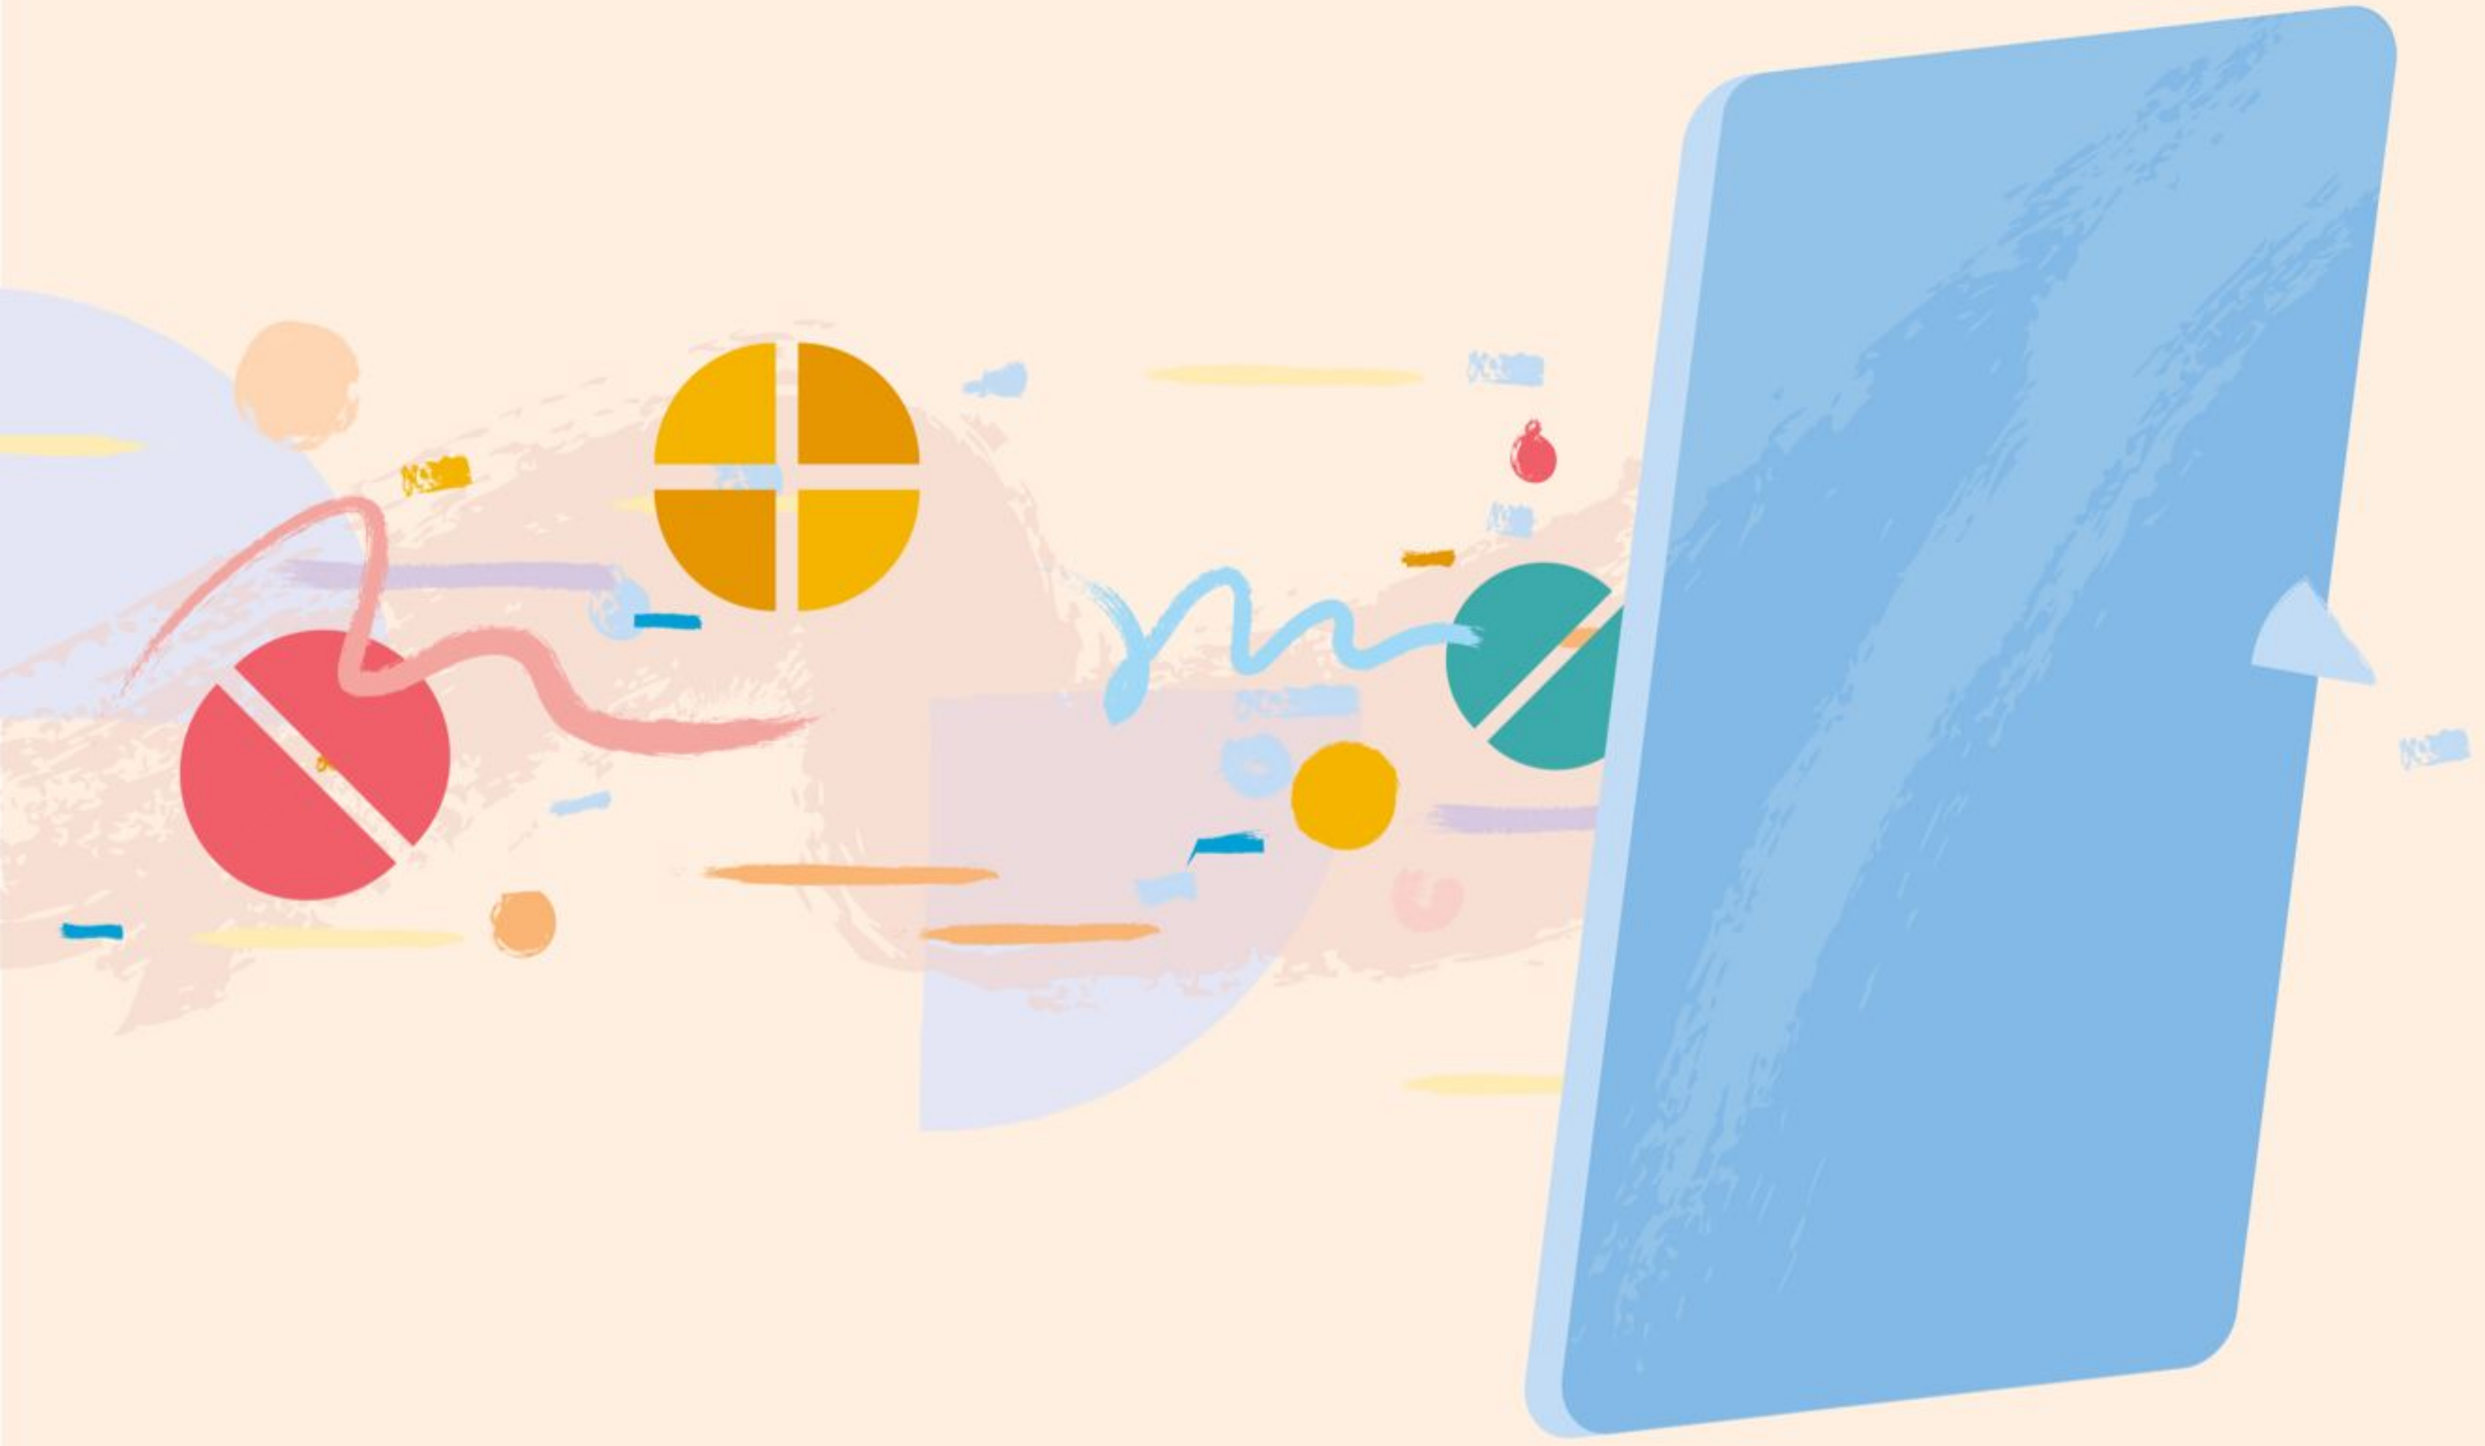

02 Product Market Fit

Market

- Private Lift went to Beta in July 2021
- Target market was ~175 sophisticated Lift advertisers.
- Good overlap with our predicted market for Private Computation

Learning #1: Advertisers are interested

> "I see this as the inevitable future of all of our ad platforms."
>
> PRIVATE COMPUTATION ADVERTISER

Learning #2: Advertisers do not feel a sense of urgency

> “If it’s better, we’d do it immediately. Better meaning campaign performance is better / [coverage is] better as a result.”

PRIVATE COMPUTATION ADVERTISER

Learning #3: More education is needed

Advertiser’s are considering MPC for the first time and need time to reason about the product constraints and privacy guarantees.

03 MPC at Scale

Lift

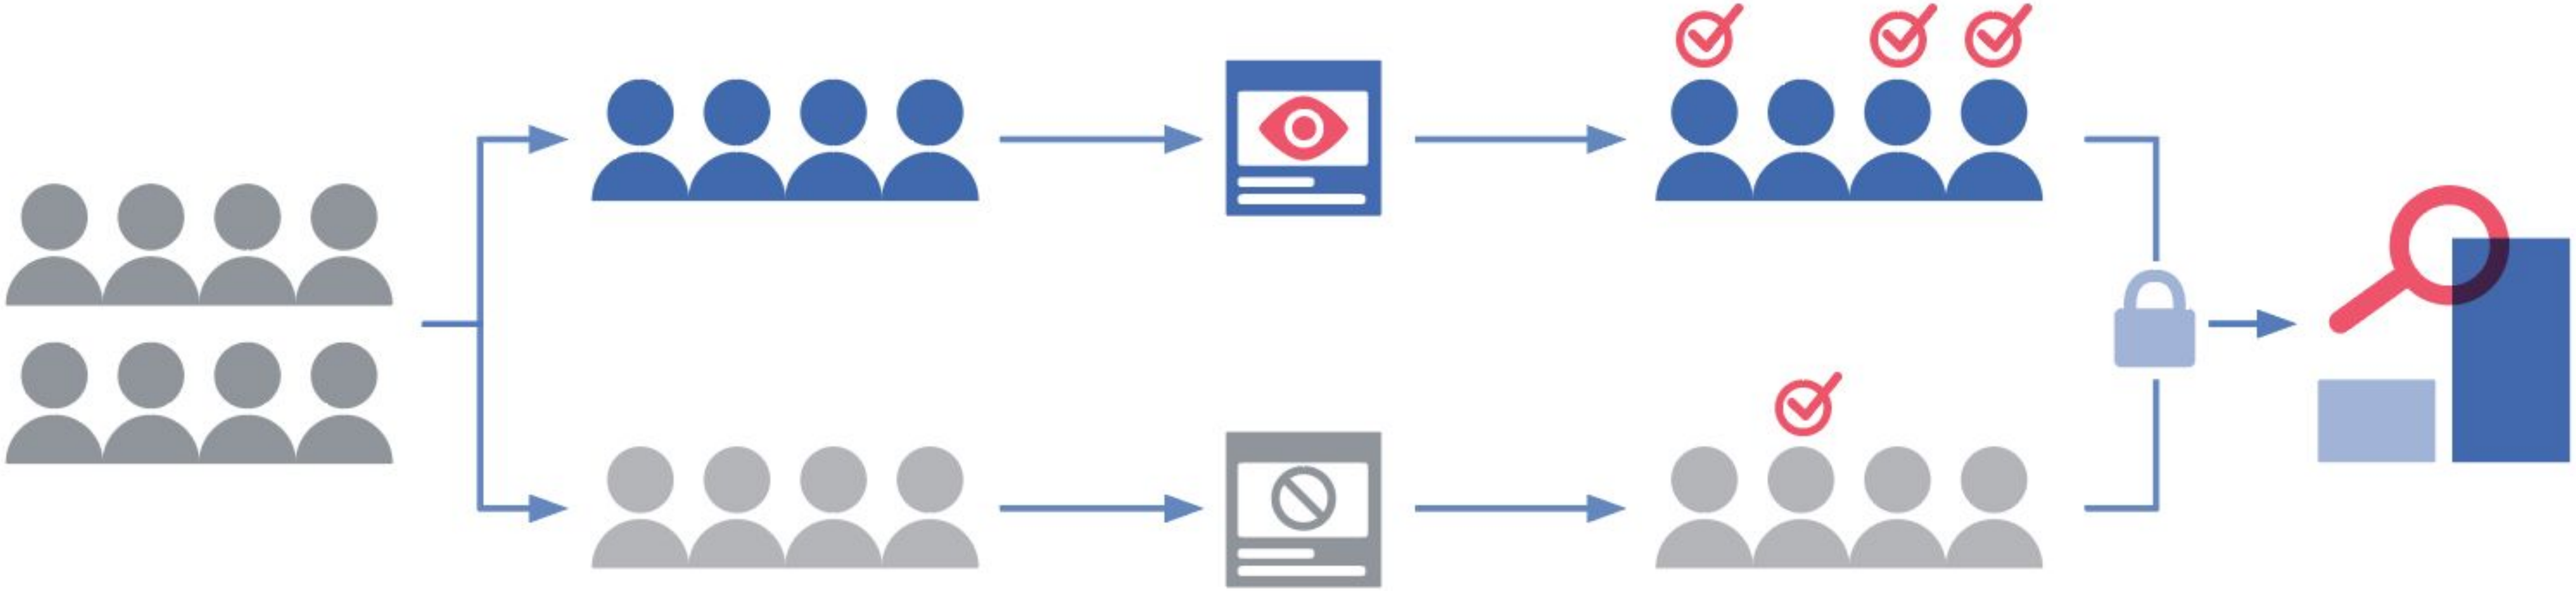

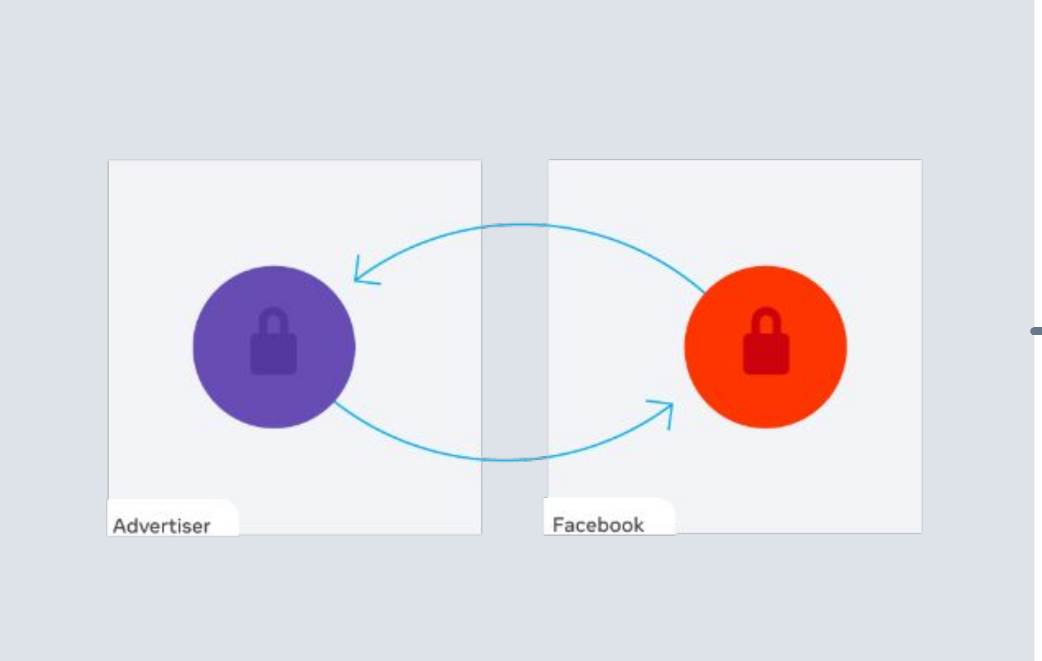

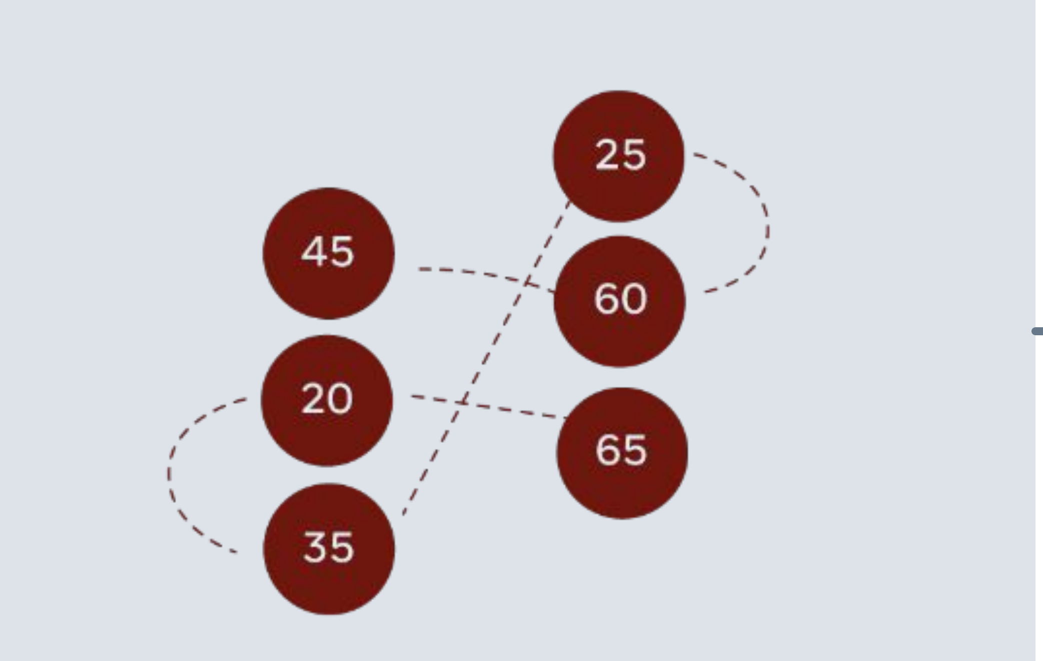

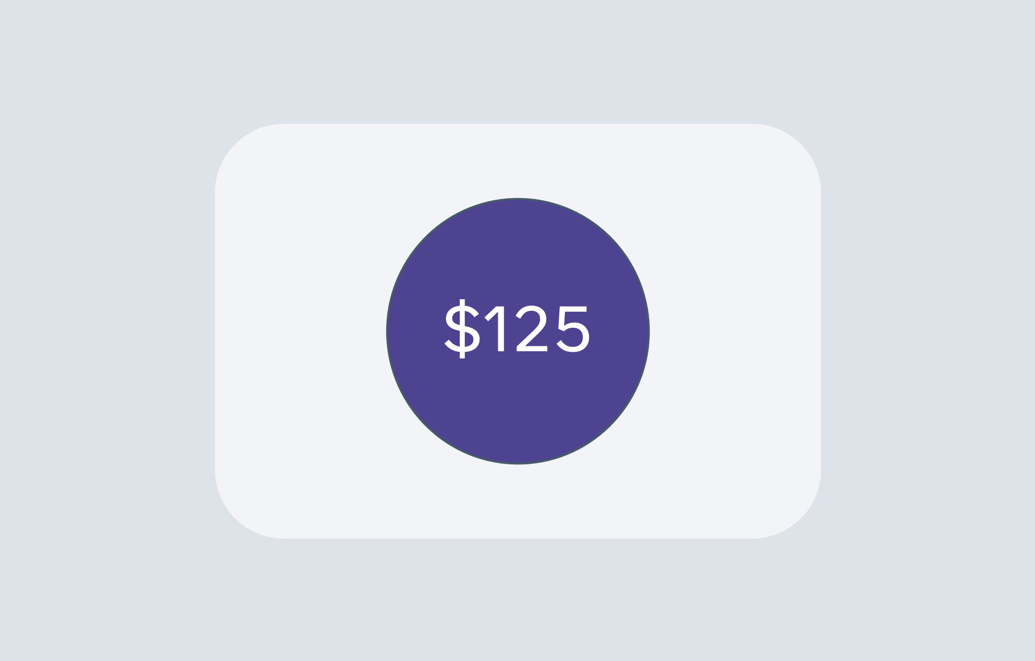

Private Record Linkage

Private records are generated from records held by the Advertiser and Meta

Private Attribution

Records have attribution rules applied to them using MPC

Private Aggregation

Attributed records are aggregated and anonymized

Private Record Linkage is Highly Complex

Private Record Linkage Challenges

Multiple Records. Both the advertiser and advertising platform typically have multiple records that map to the same unit you are trying to measure.

Algorithm must support sharding. In a multi-record protocol, the multi-identifiers cannot be confined to a single shard and a cross shard communication is required. This is highly complex and performance intensive.

Custom protocols are expensive to build, debug, and maintain. Each protocol variant can take up to a year from research to production-ready.

Quality of Records. The quality of records varies significantly across advertisers and degrade rapidly.

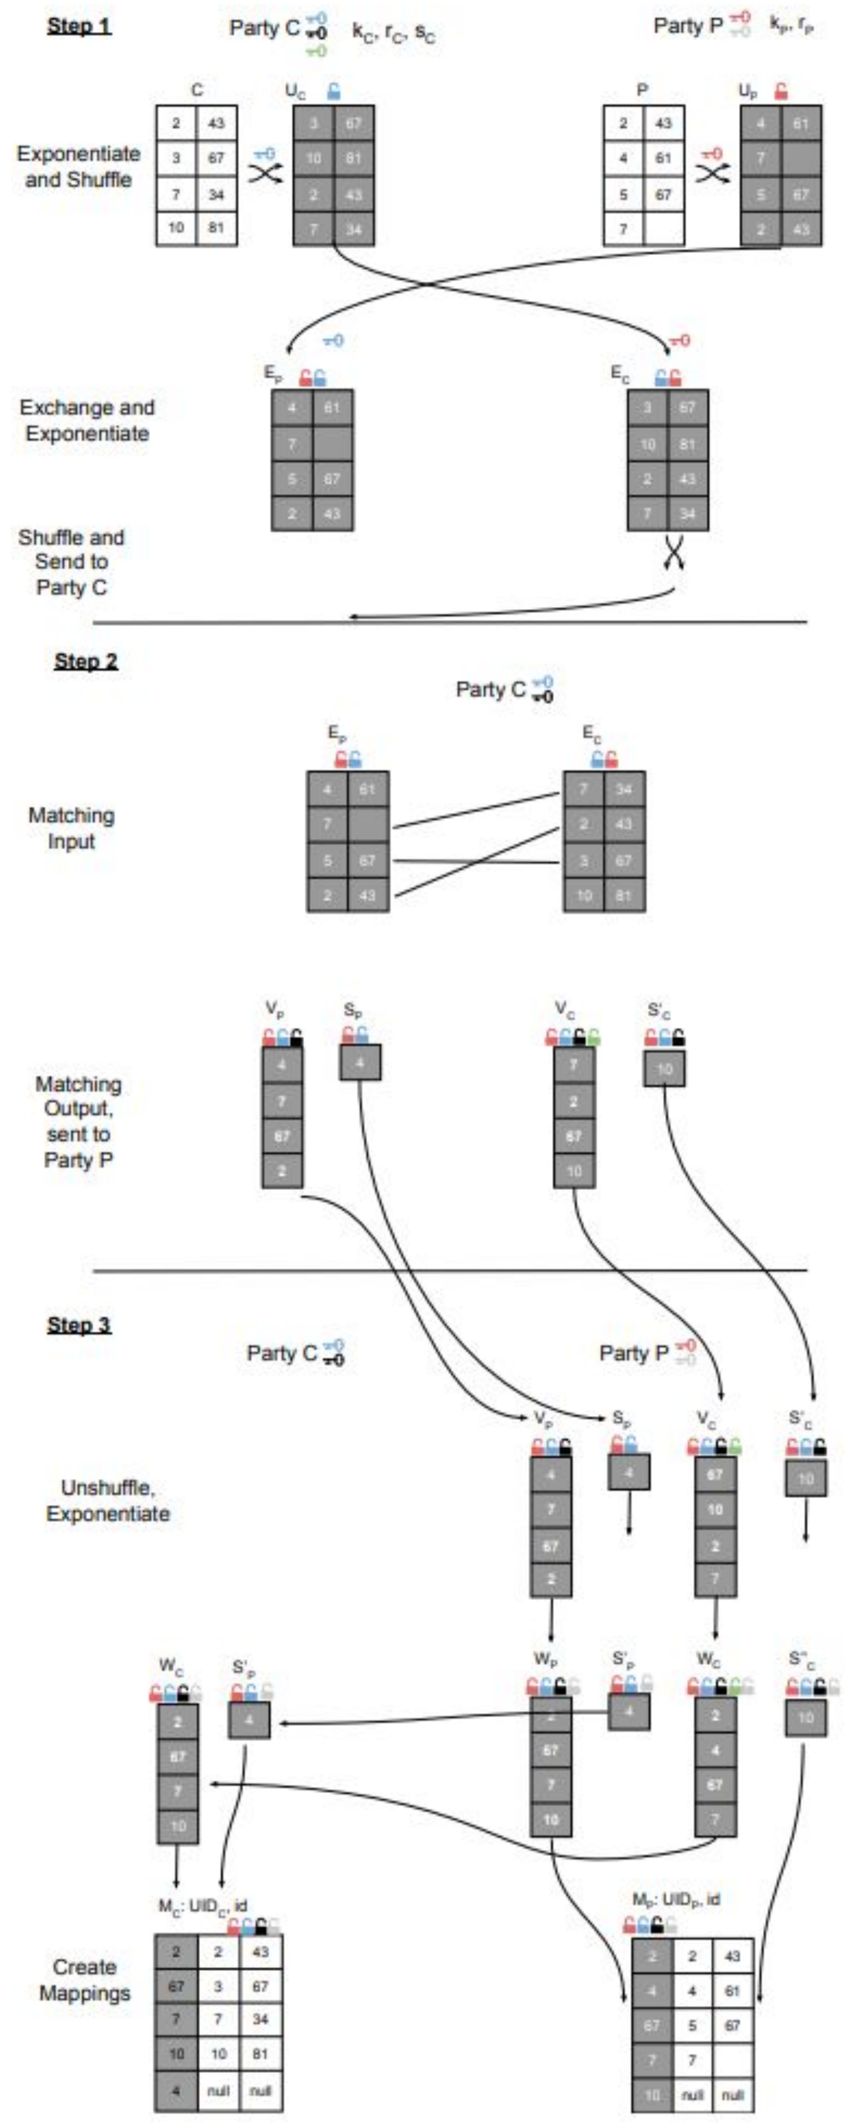

Supporting large computations

1B

Target maximum input rows into Private Lift

2hrs

Target maximum time to complete the computation

1/1000

Target computation cost relative to campaign spend

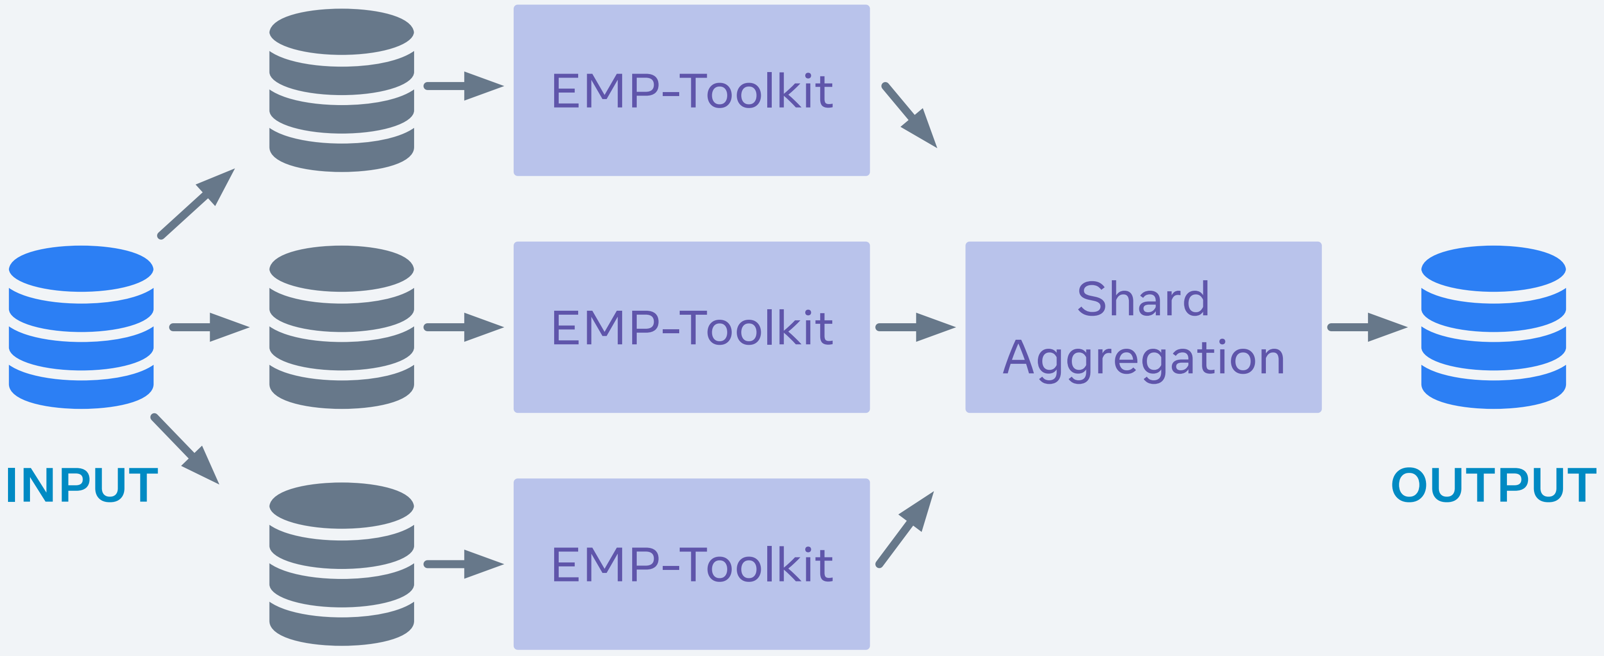

Maintaining input privacy while scaling horizontally

Cost of a Private Computation

Advertisers typically allocate a budget for measuring their ads. We are well within that.

Typical cost of a Private Attribution computation is comparable to cost of a movie ticket.

Cost of Adding Features

Despite the algorithm being simpler, Private Lift is a much more full featured product than Private Attribution. Supporting the full suite of Private Lift features is costly.

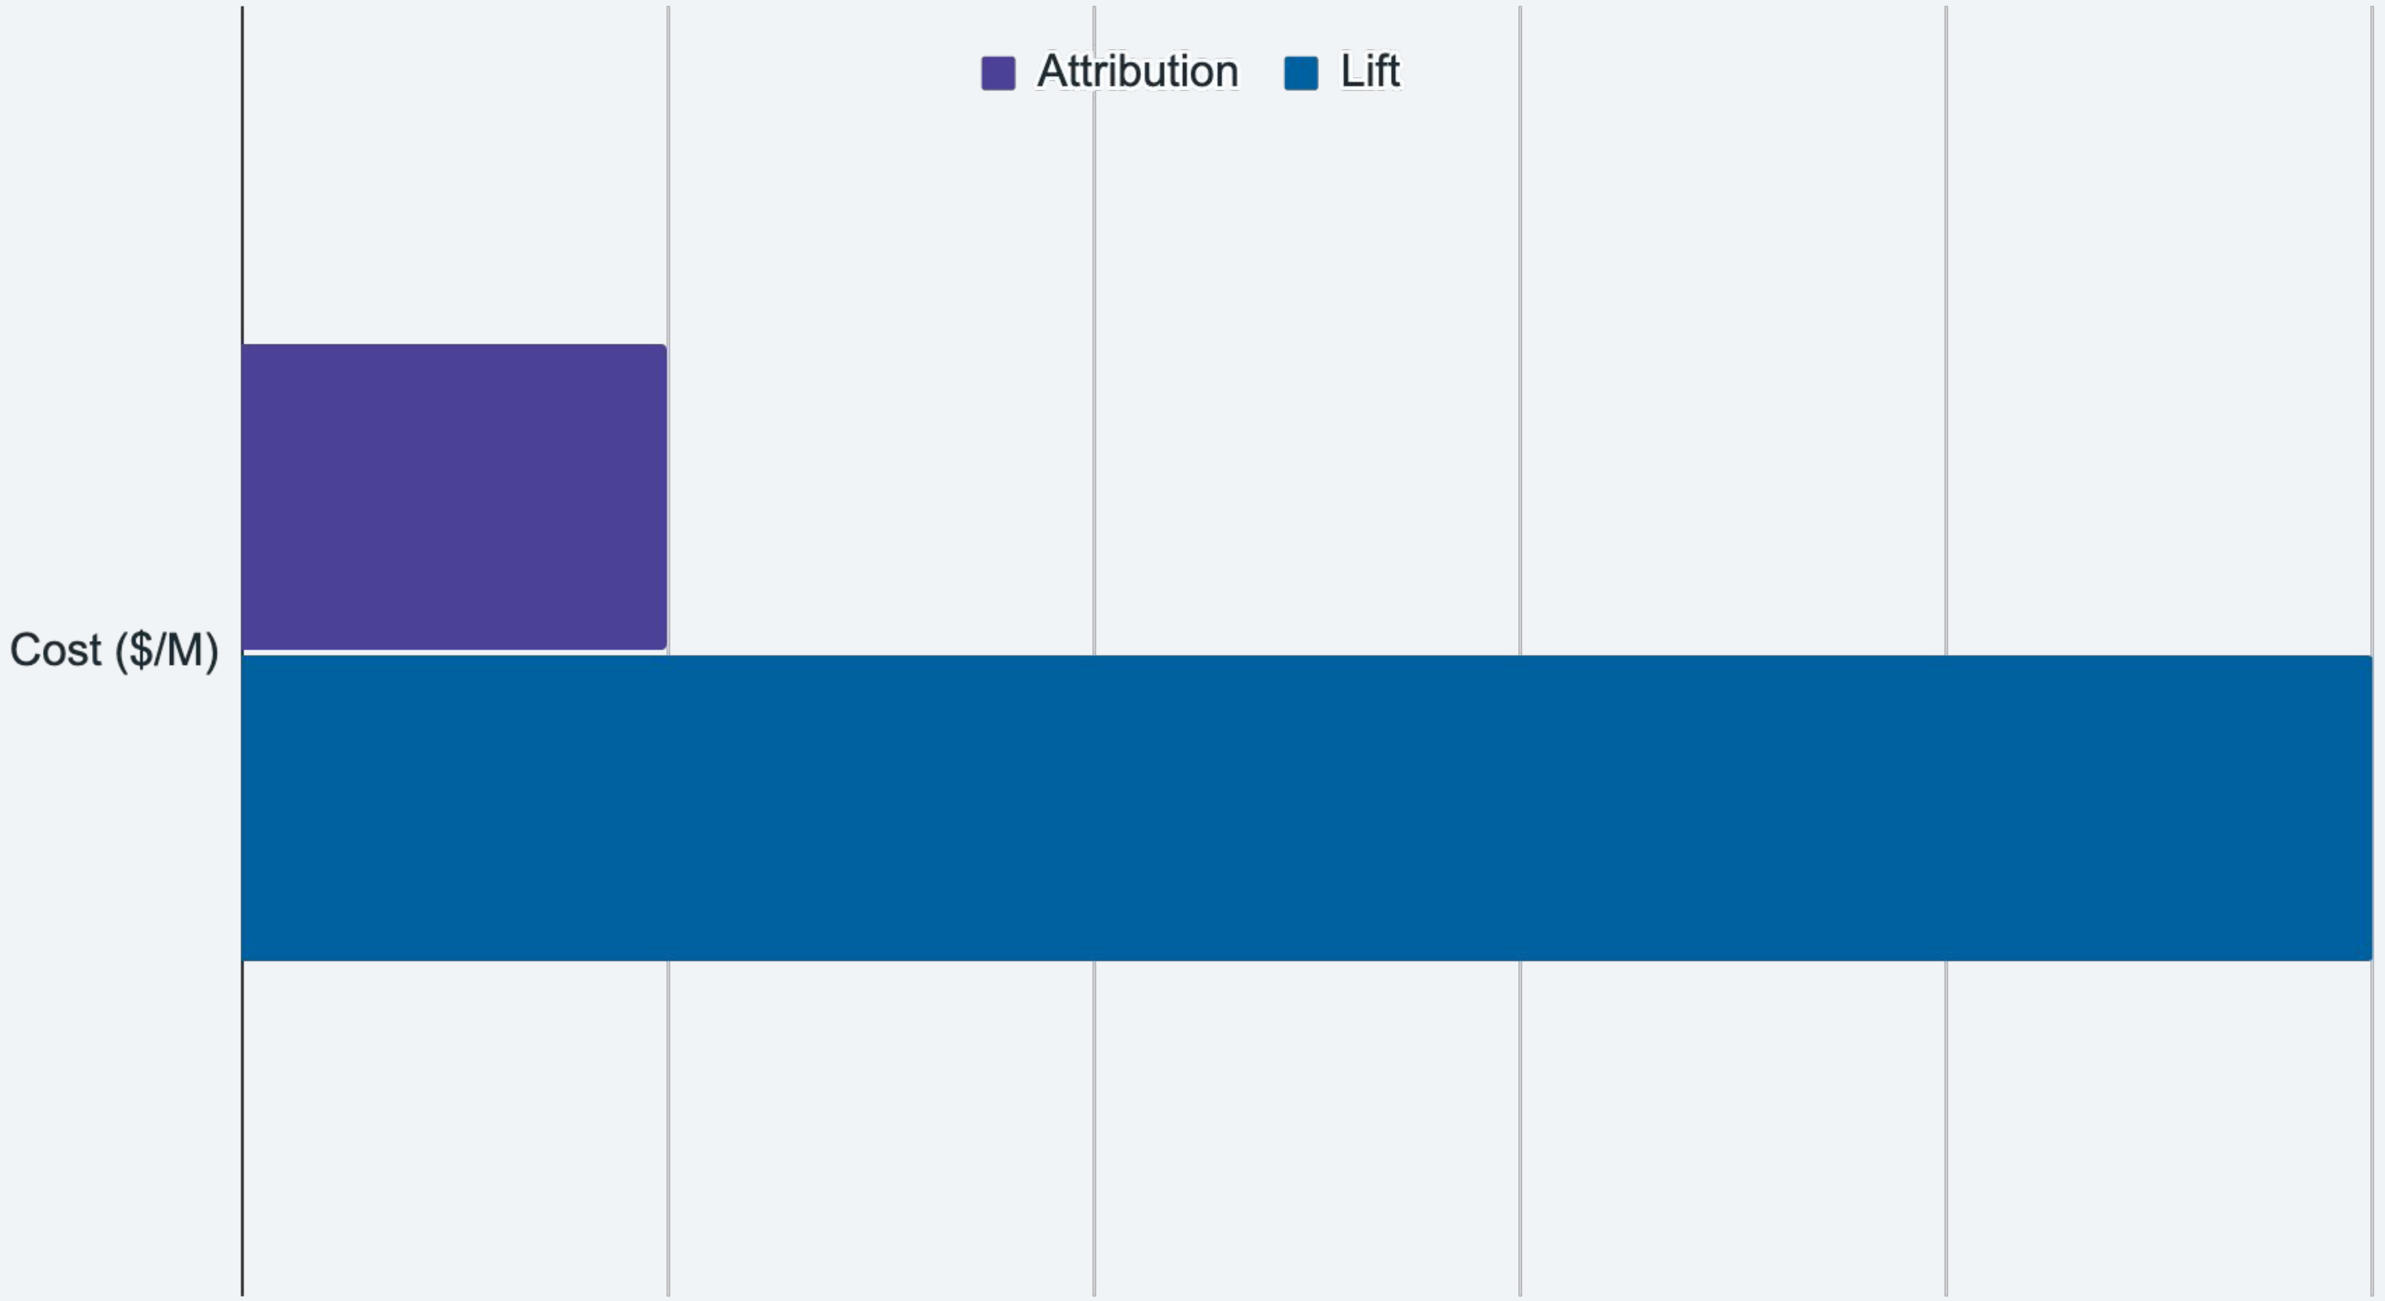

Private Computation Framework

New MPC Engine from Meta. V2 switches from EMP-Toolkit backend to a Meta-built MPC Engine.

Fast and Efficient. 100x reduction in network usage and 4x improvement in performance.

Open Sourced. Fully production-ready with a whitepaper this half. Already open-sourced. We welcome any contributions!

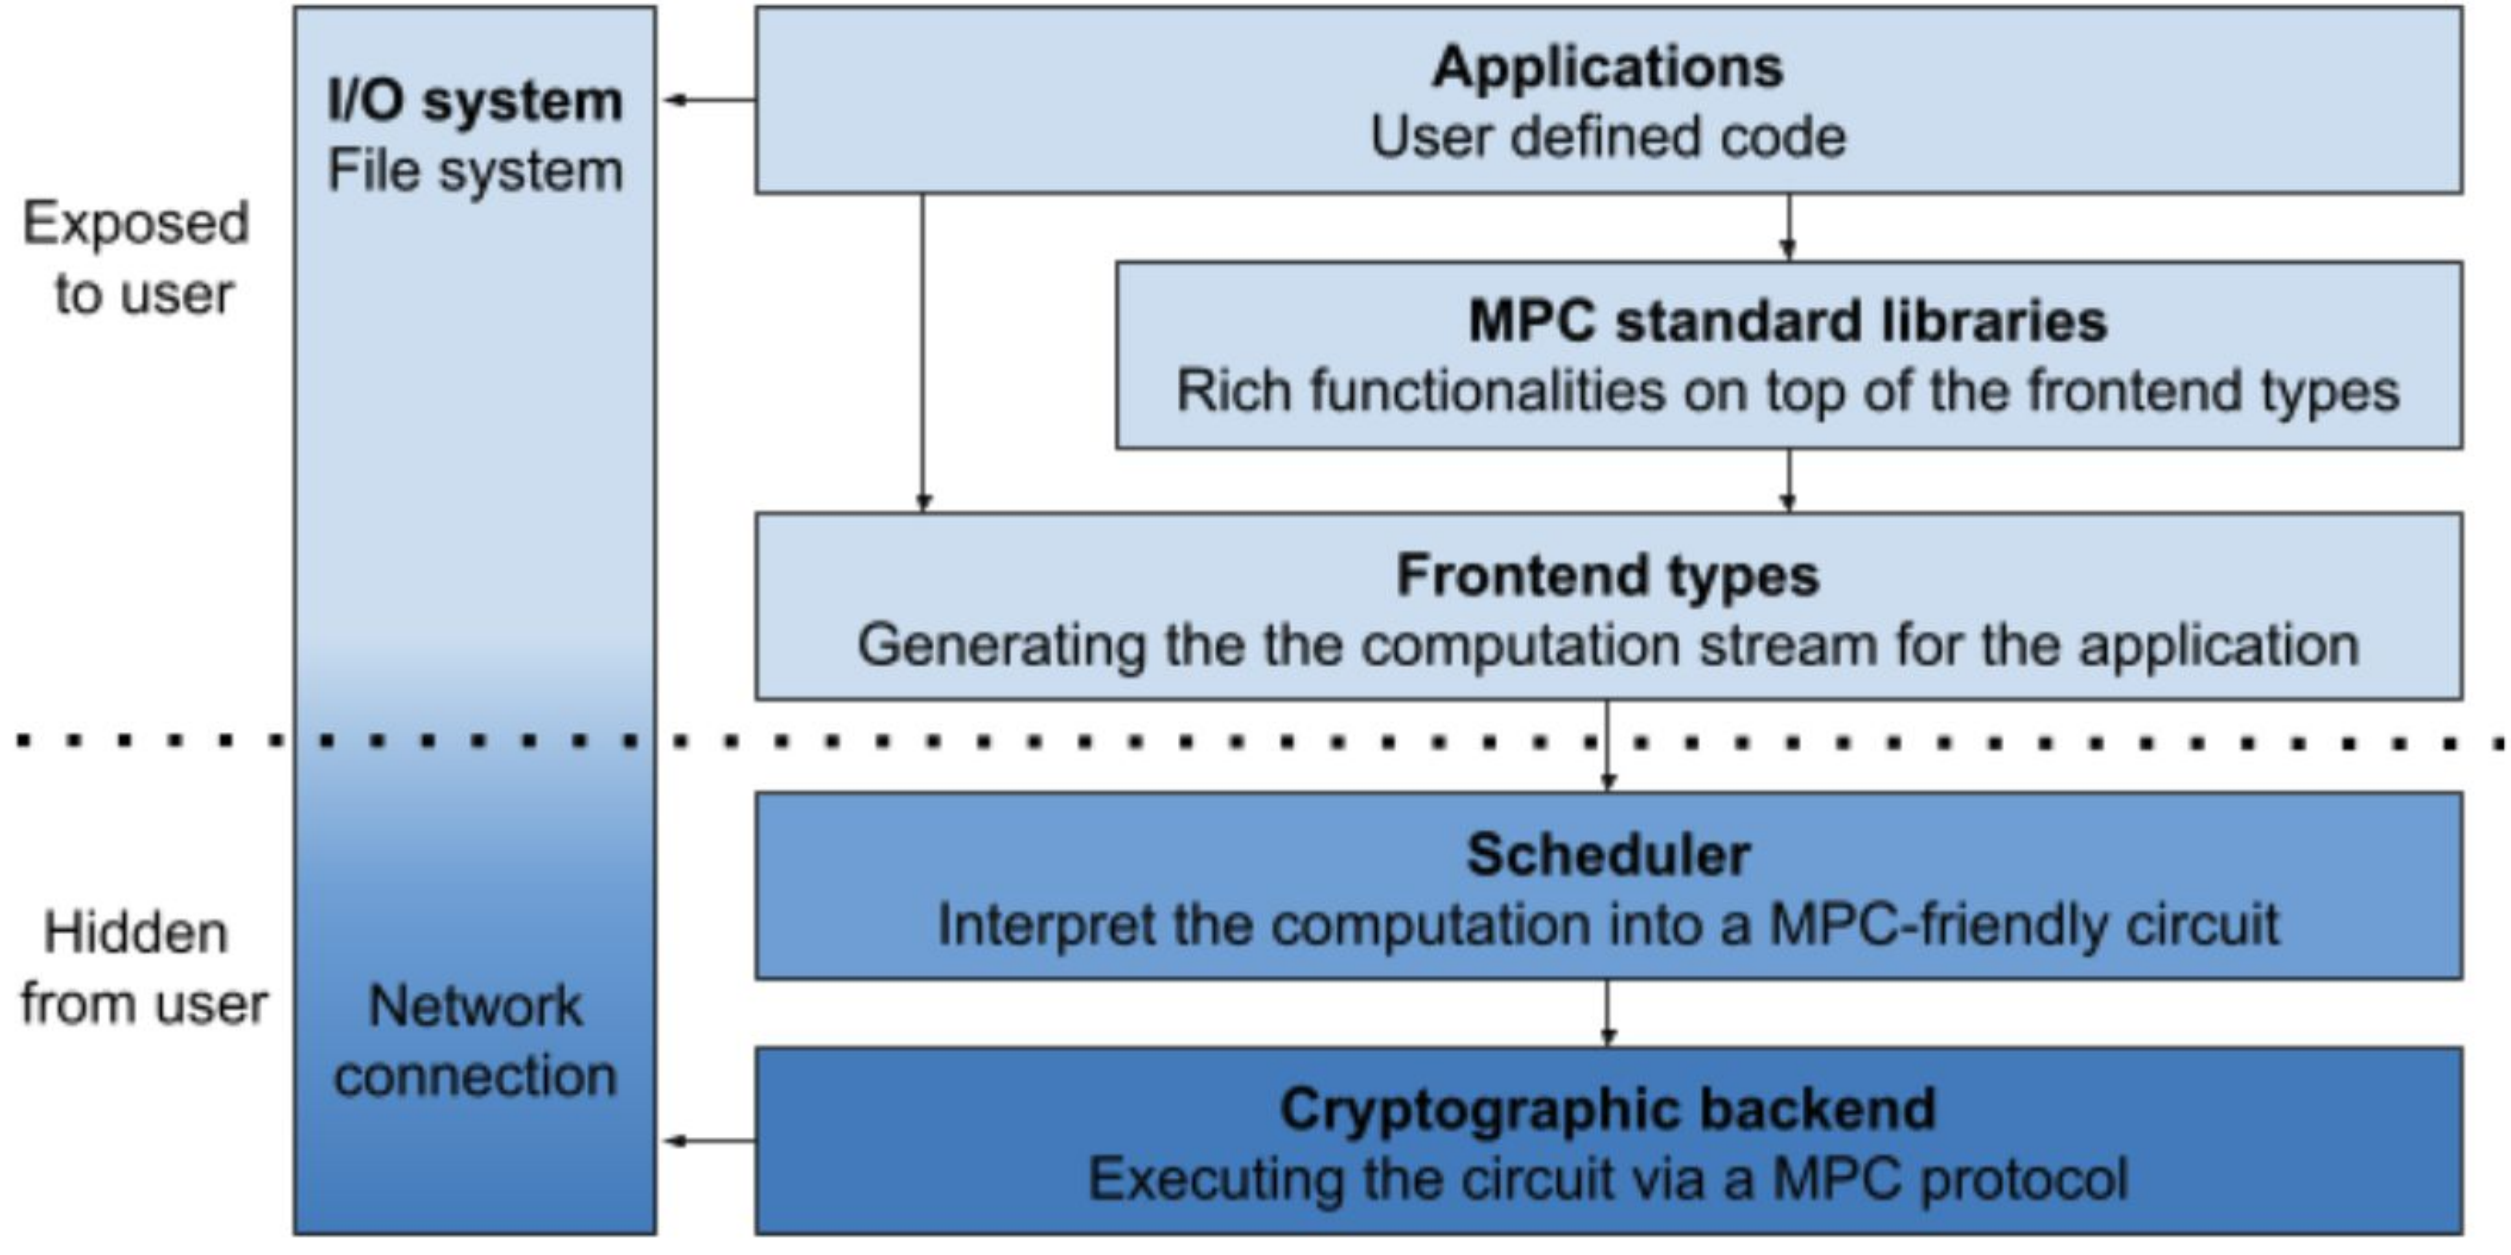
Exposed to user
Hidden from user
I/O system
File system
Network connection
Applications
User defined code
MPC standard libraries
Rich functionalities on top of the frontend types
Frontend types
Generating the the computation stream for the application
Scheduler
Interpret the computation into a MPC-friendly circuit
Cryptographic backend
Executing the circuit via a MPC protocol

fb.me/pcf

Supporting a large number of advertisers

10m+

Meta advertisers

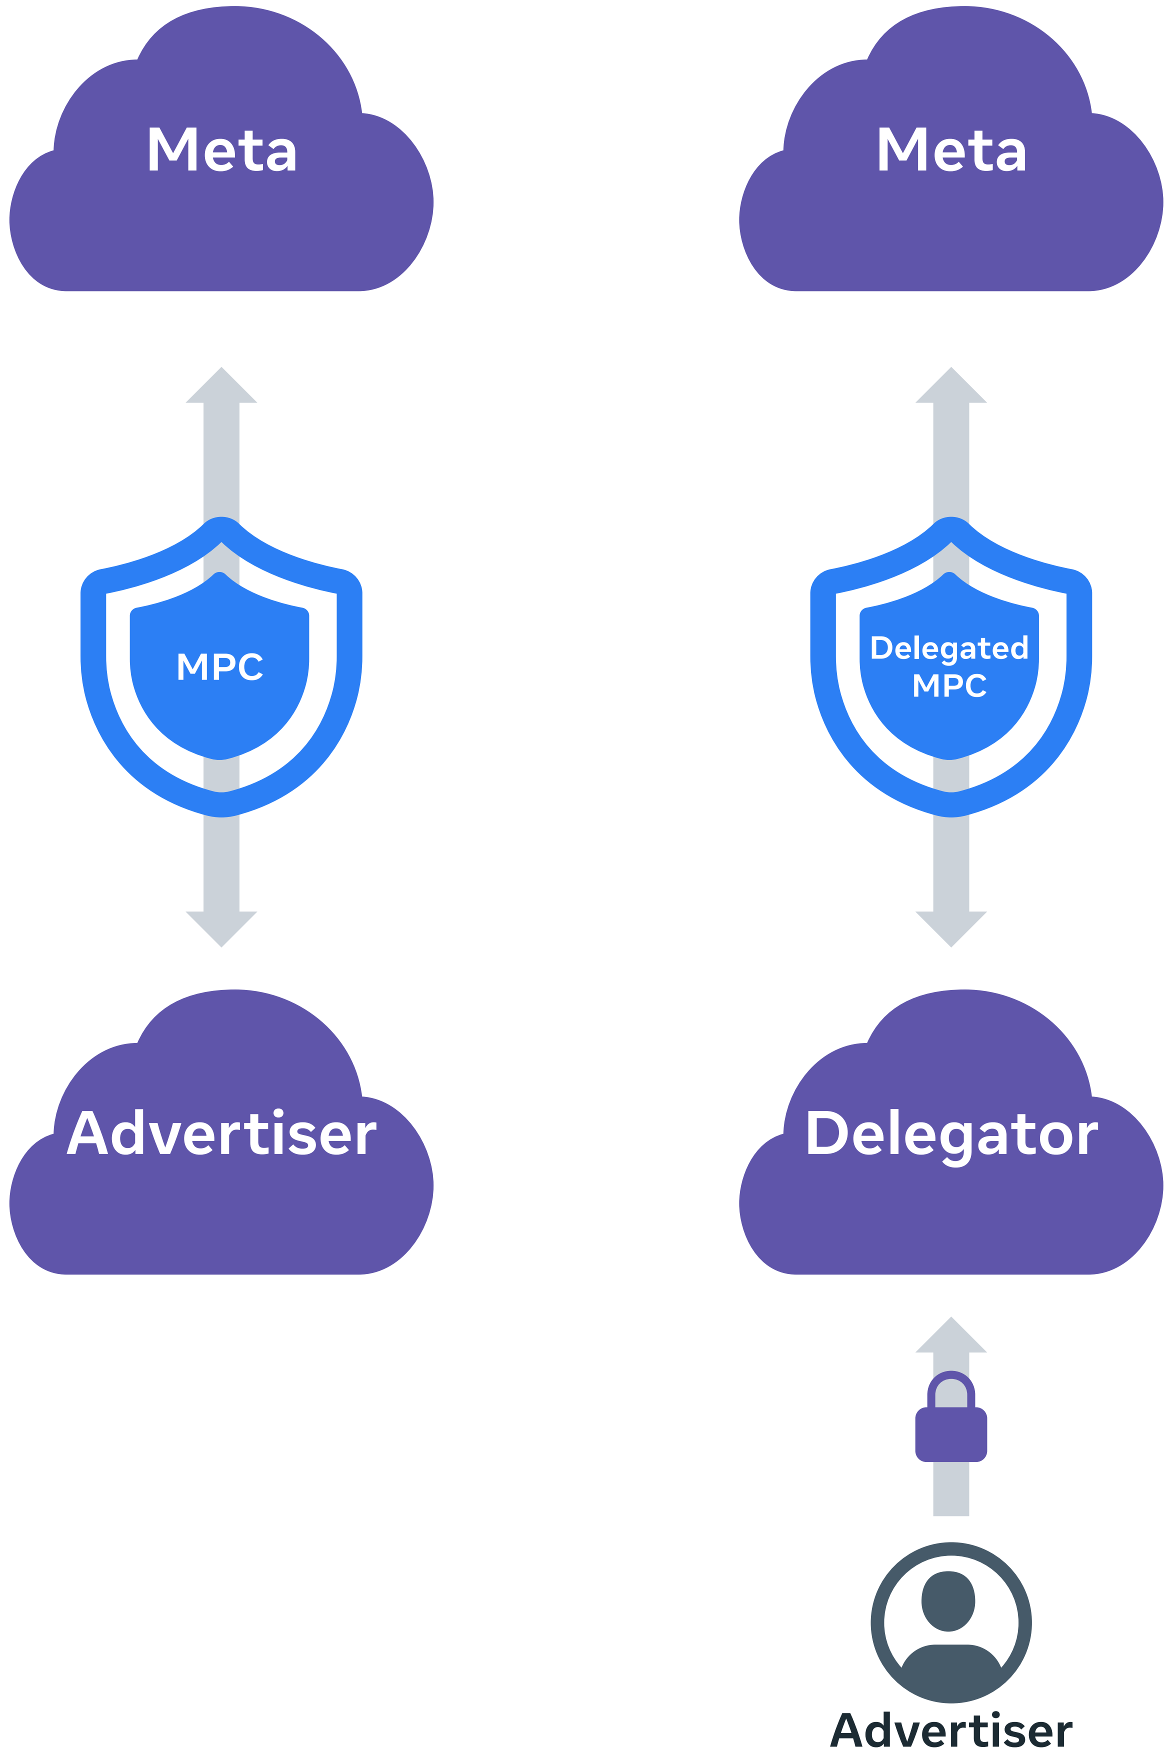

Deploying MPC is difficult

Months

Average time for an advertiser to onboard to Private Computation today

Hours

By end of 2022

Infrastructure

Symmetric. Advertiser deploys the same code that Meta runs.

Serverless. Containers are spun up for computation that communicate over VPC peering.

Simple Deployment. Self-service, no-code set up of your infrastructure, with built-in validation.

Automated Data Management. Robust data pipeline that can optionally leverage your existing data setup with Meta.

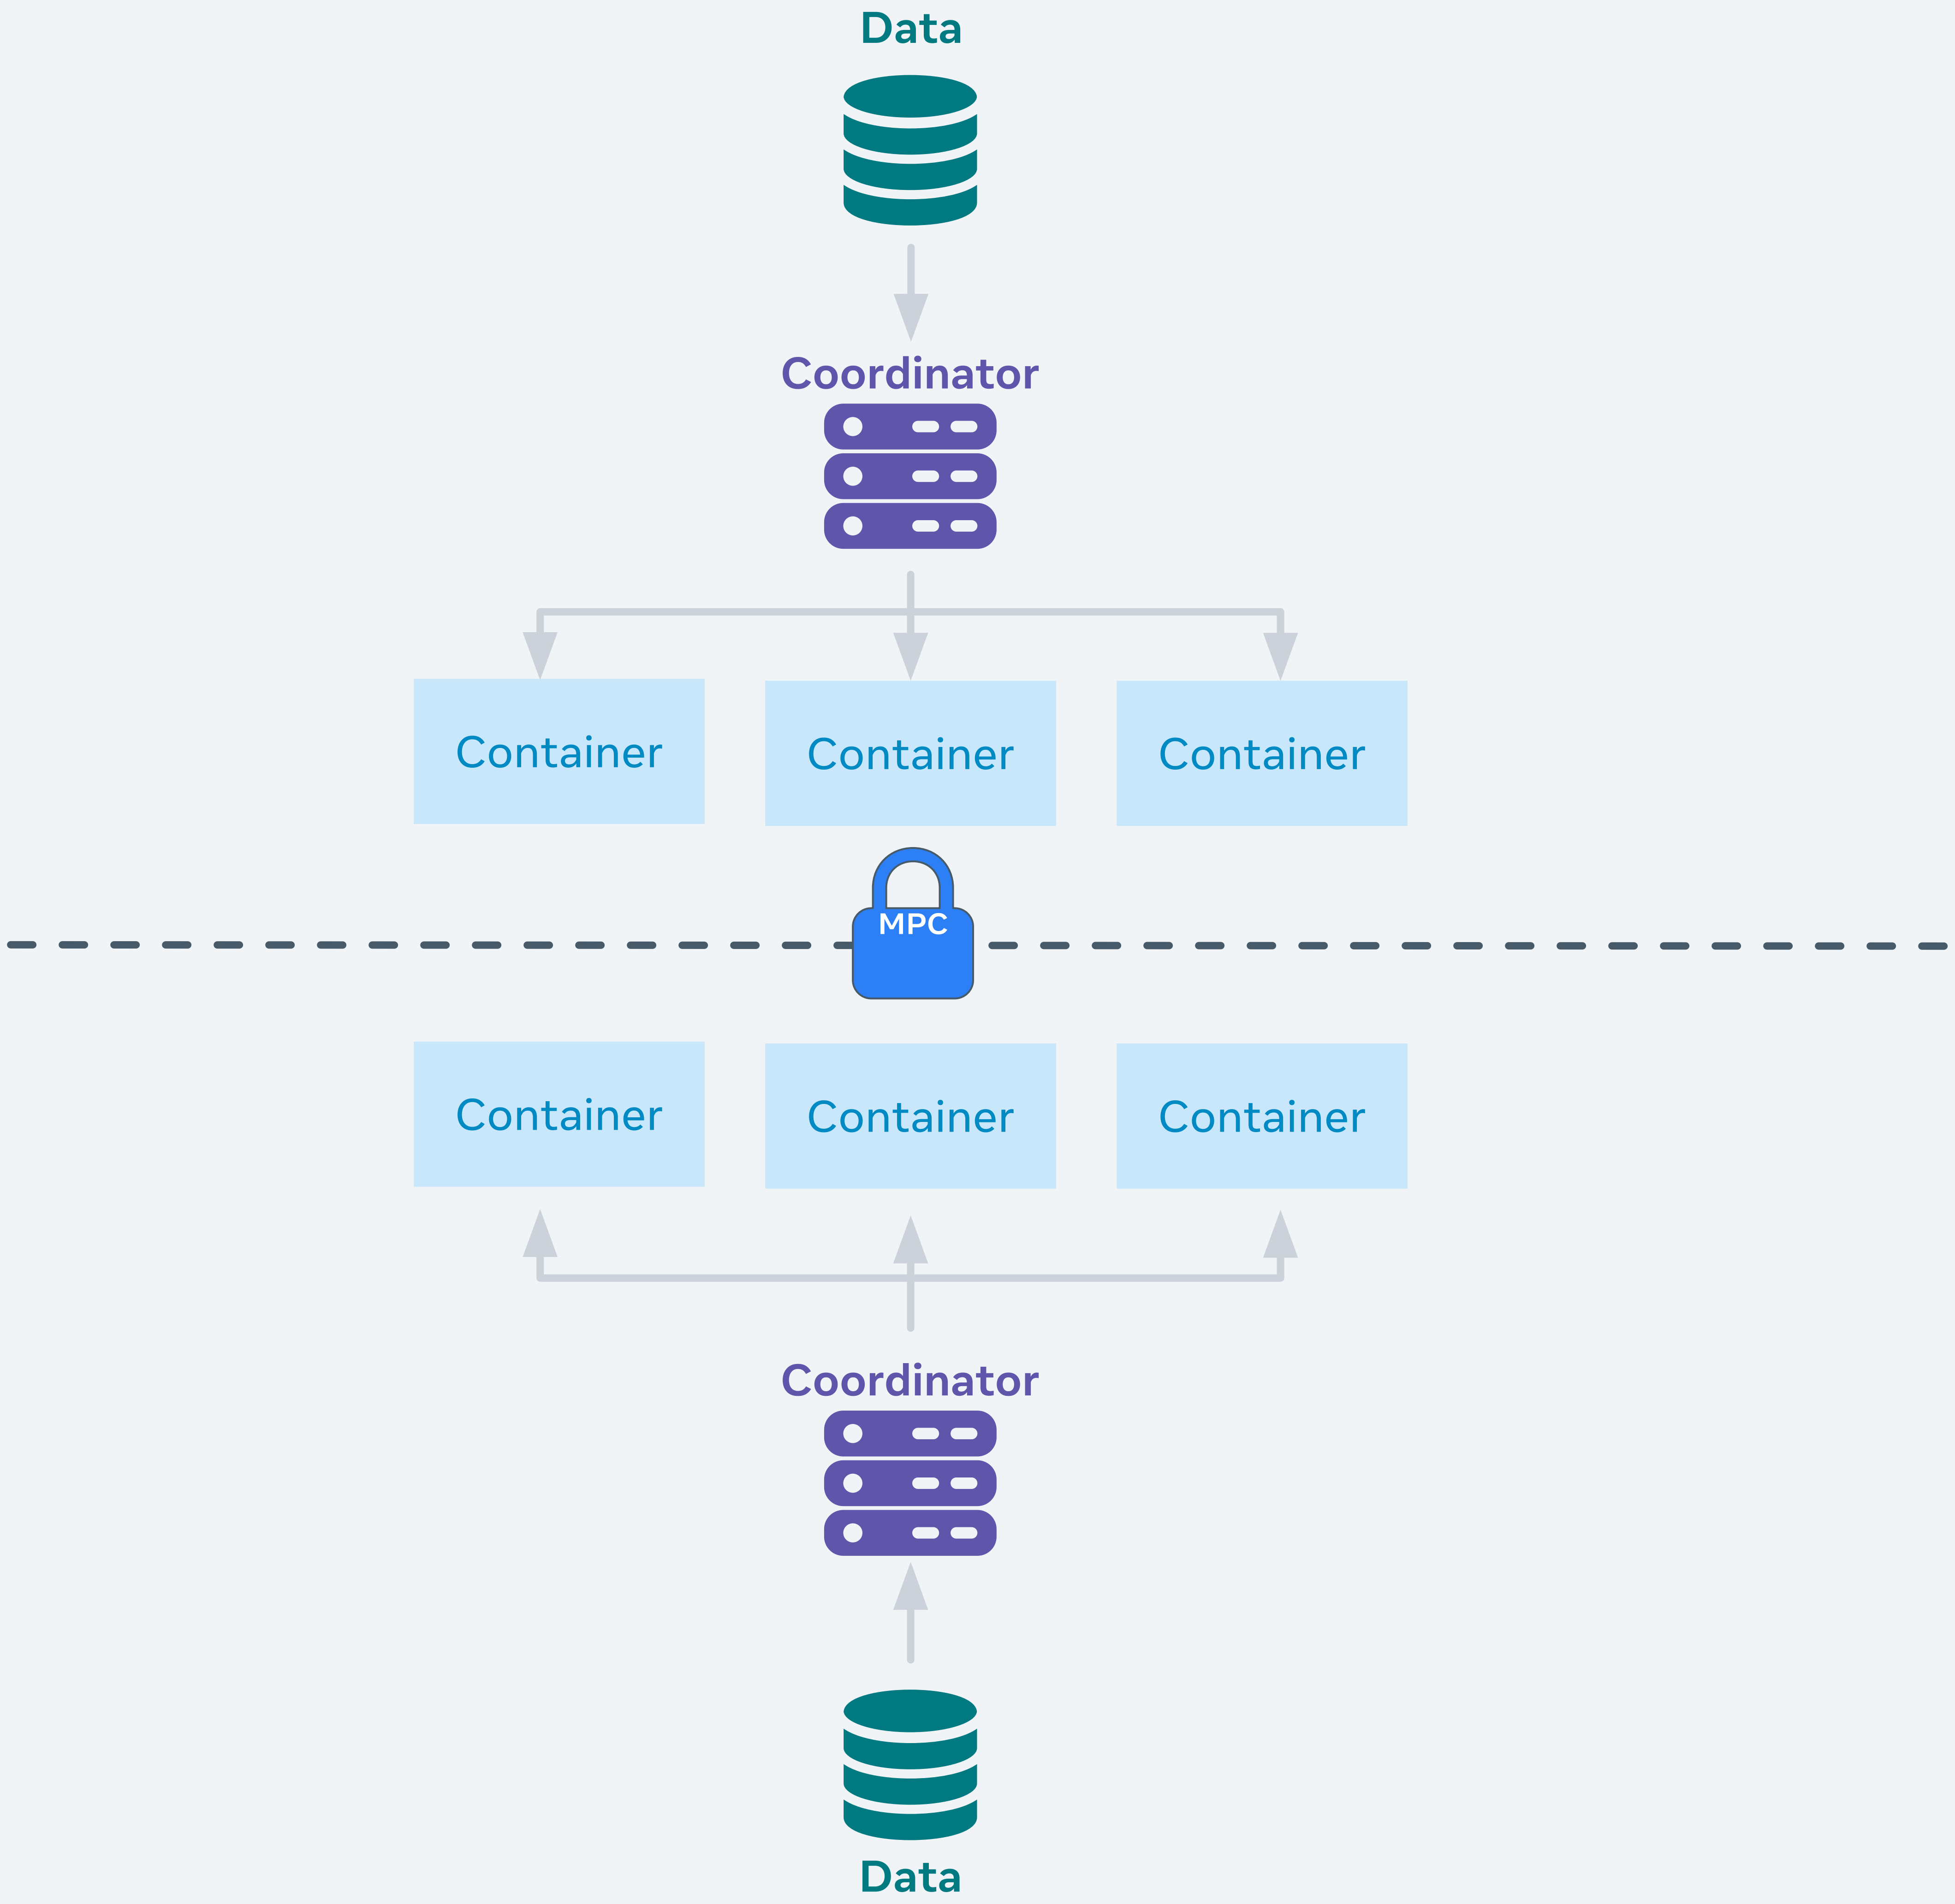

fb.me/pcp

Delegated Integrations

MPC Consortium

Smaller advertisers could leverage a central consortium to perform their computation, without trusting any of the members.

Partner Integrations

Advertisers often already delegate their commerce or measurement needs to a third-party service.

Managed MPC

We can build secure systems that allow other parties to manage advertiser’s MPC infrastructure.

05 Closing

Private Computation is Open Source

Lift and Attribution: [fb.me/pcs](fb.me/pcs)

Platform: [fb.me/pcp](fb.me/pcp)

MPC Engine: [fb.me/pcf](fb.me/pcf)

Record Linkage: [fb.me/pid](fb.me/pid)

Privacy Tech at Meta

Data for Good
Private Data for Research

Metaverse
Permissioned Distributed Ledger Tech

Advertising
Federated Learning

Virtual Reality
On-Device Processing

Messaging
End-to-End Encryption

Authentication
Anonymous Credentials

Analytics
Aggregation & Anonymization

Infrastructure
Data Security

Commerce / Financial Tech
Blockchain

Advertising
Private Computation

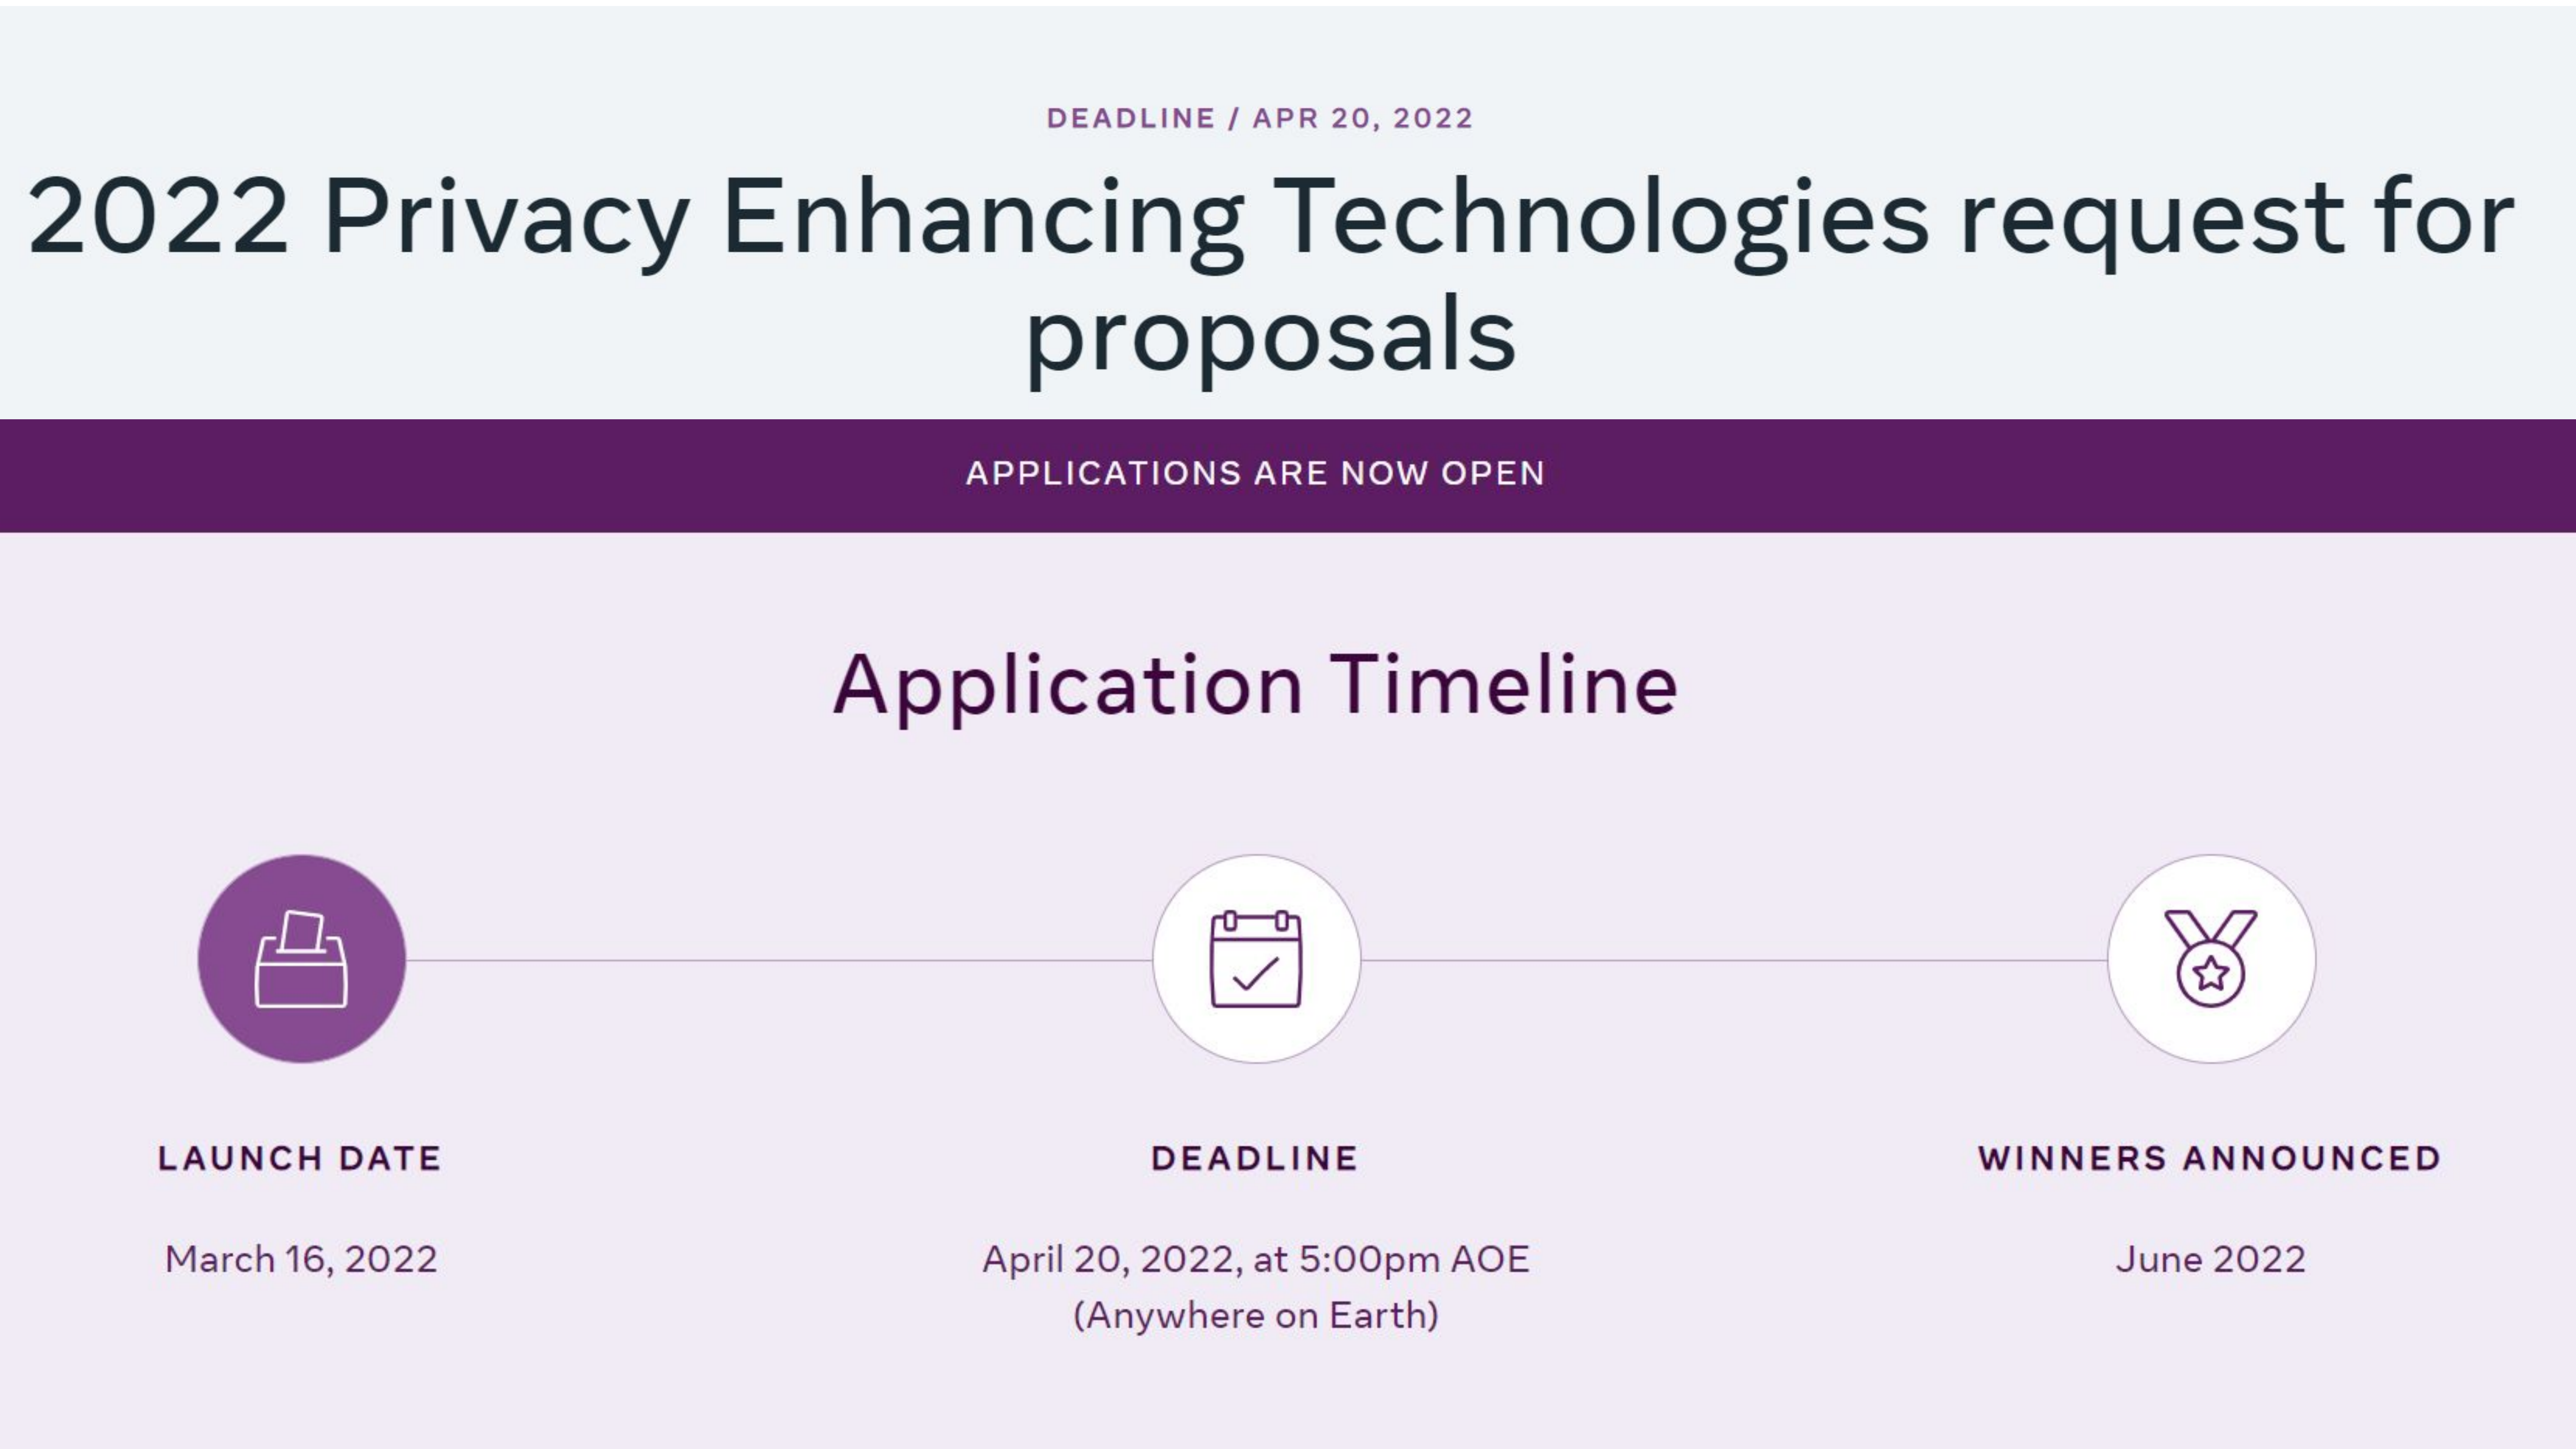

DEADLINE / APR 20, 2022

2022 Privacy Enhancing Technologies request for proposals

APPLICATIONS ARE NOW OPEN

Application Timeline

LAUNCH DATE	DEADLINE	WINNERS ANNOUNCED
March 16, 2022	April 20, 2022, at 5:00pm AOE (Anywhere on Earth)	June 2022

James Reyes

Private Computation @ Meta
jlreyes@fb.com

Sanjay Saravanan

Cryptography Research @ Meta
ssanjay@fb.com

Q&A

Meta