Partitioning as an Architecture for Privacy

draft-kpw-iab-privacy-partitioning

Tommy Pauly, Mirja Kühlewind, Chris Wood
IAB Open
IETF 115, November 2022, London

Many newer groups in the IETF are working on improving **user privacy** by **separating data** between entities

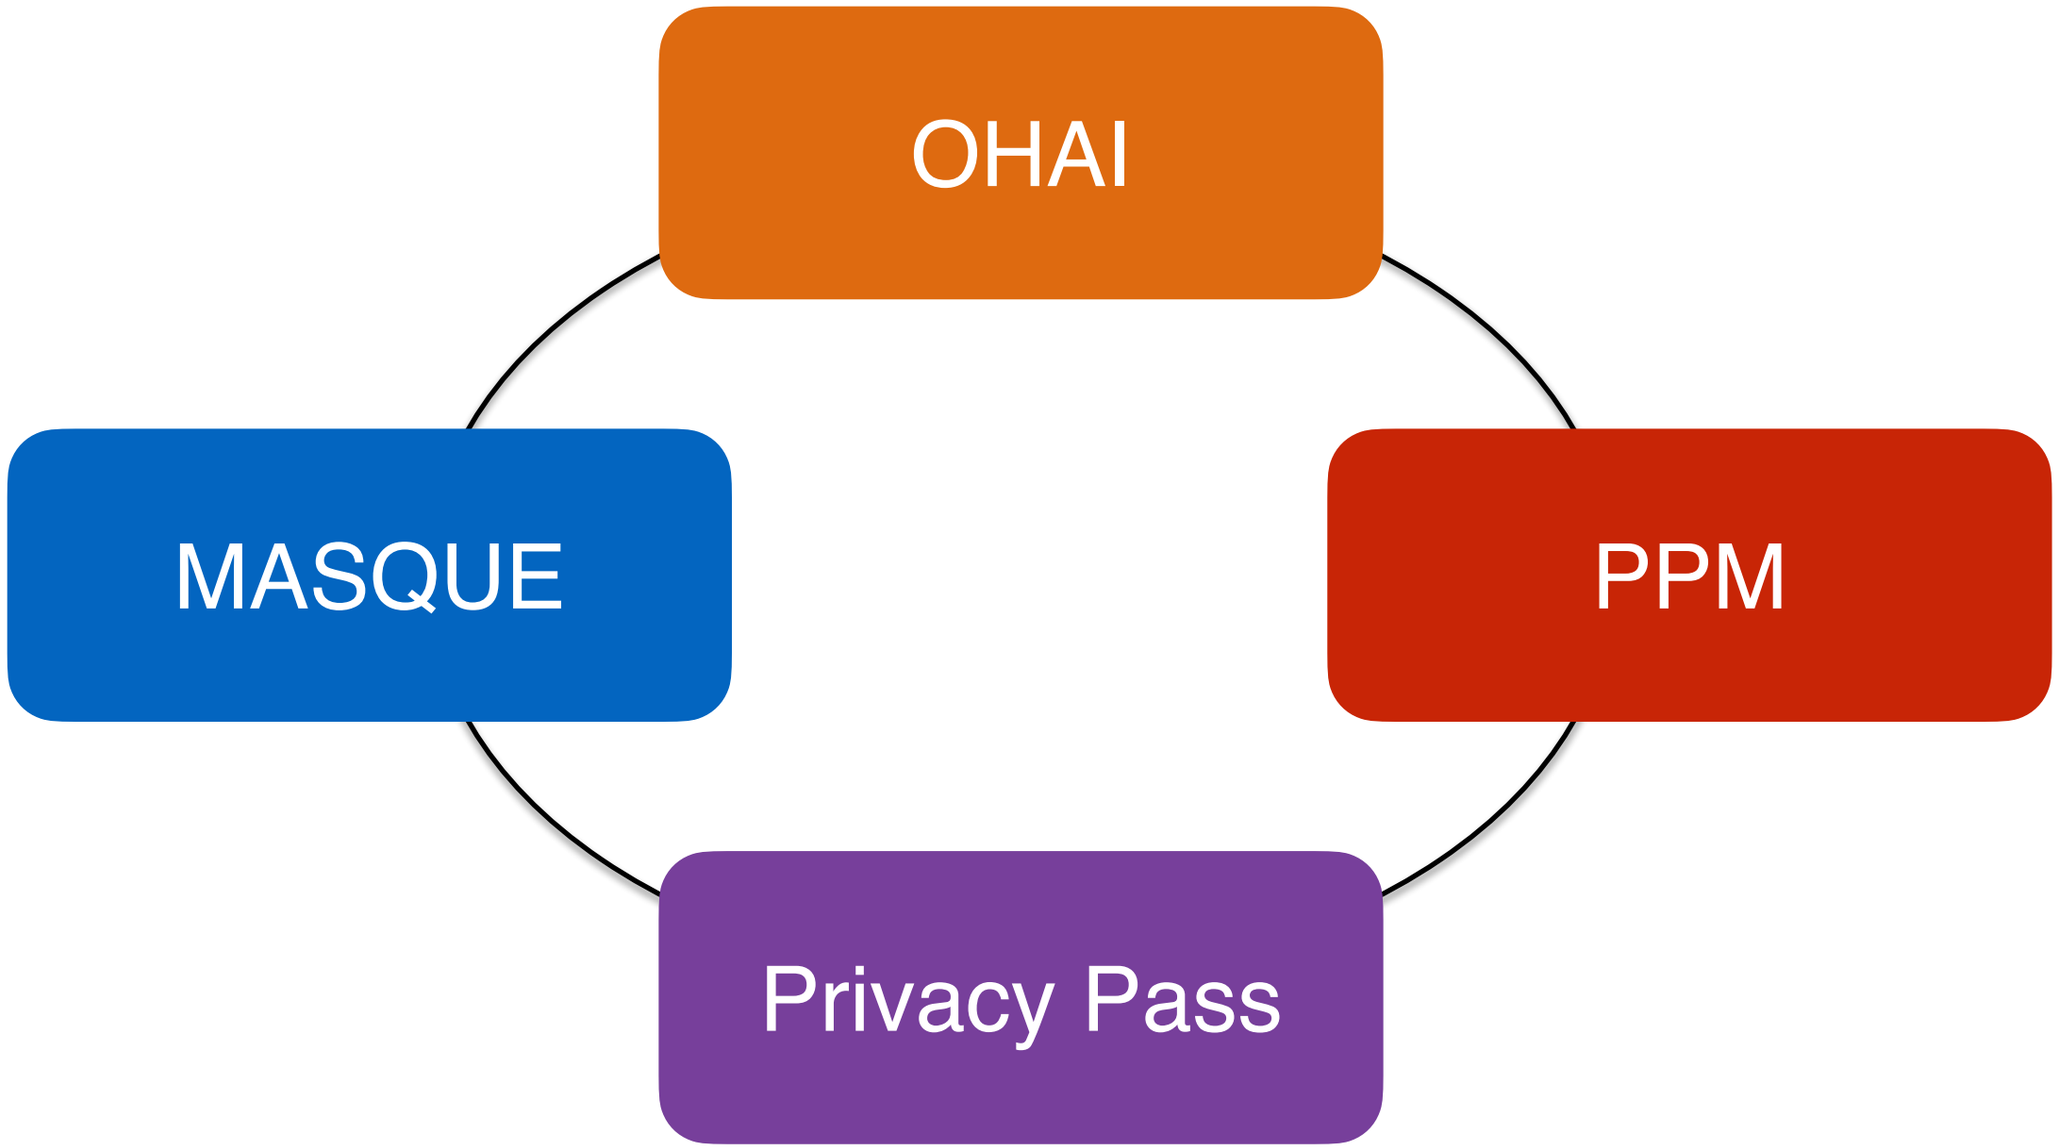
OHAI
MASQUE
PPM
Privacy Pass

Different use cases all benefit from partitioning

- Separating Client IP address from user data
- Separating user authorization from what content they access
- Separating client identity from metrics they upload

Partitioning for privacy is an emerging architecture pattern in Internet protocols

We need to work on ways to discuss, analyze, and evaluate these protocols

Privacy Contexts

A context is a group of entities that share a view of data and metadata

Within a context, data and identifiers can be trivially correlated

We identify two techniques for partitioning contexts:

1. **Encryption**, which allows partitioning along a network path or in a connection

2. **Separate connections** across time or space

Example: Oblivious HTTP

Context 1: End-to-end messages

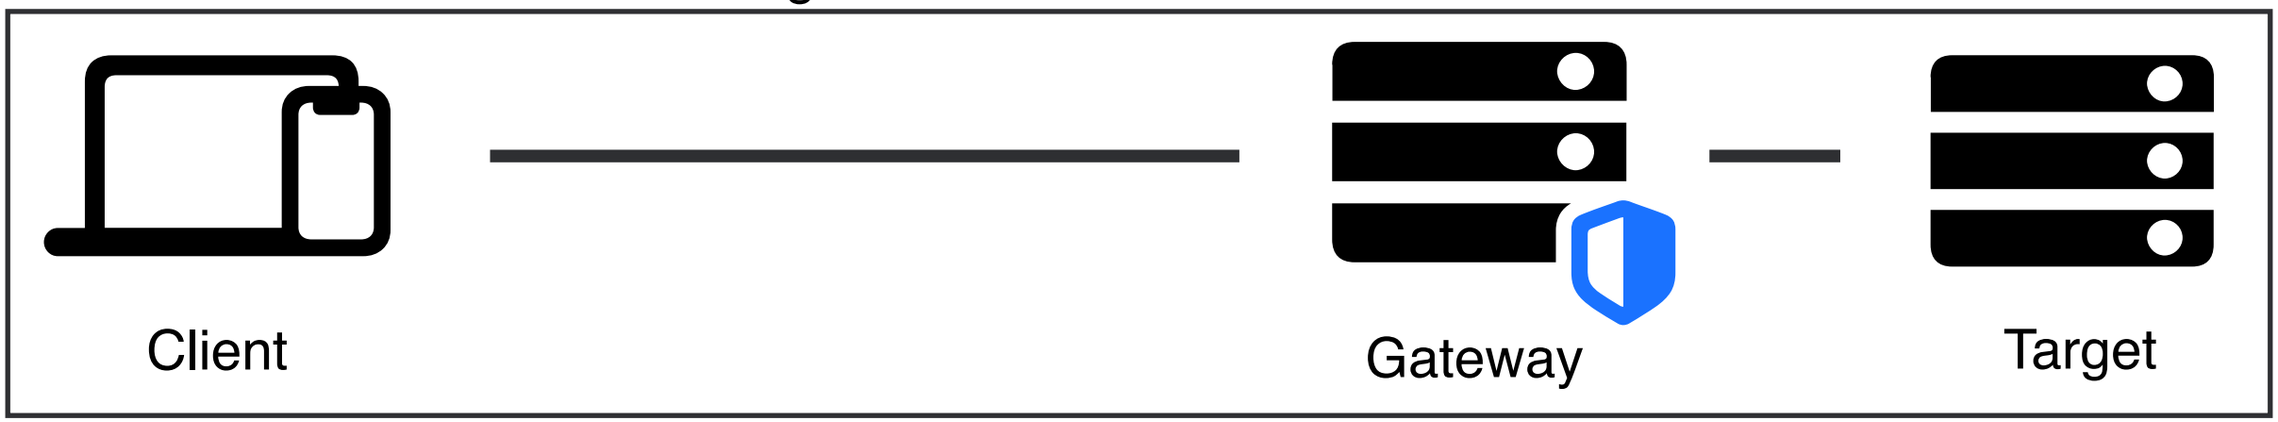

Context 2: Client-to-gateway encrypted messages

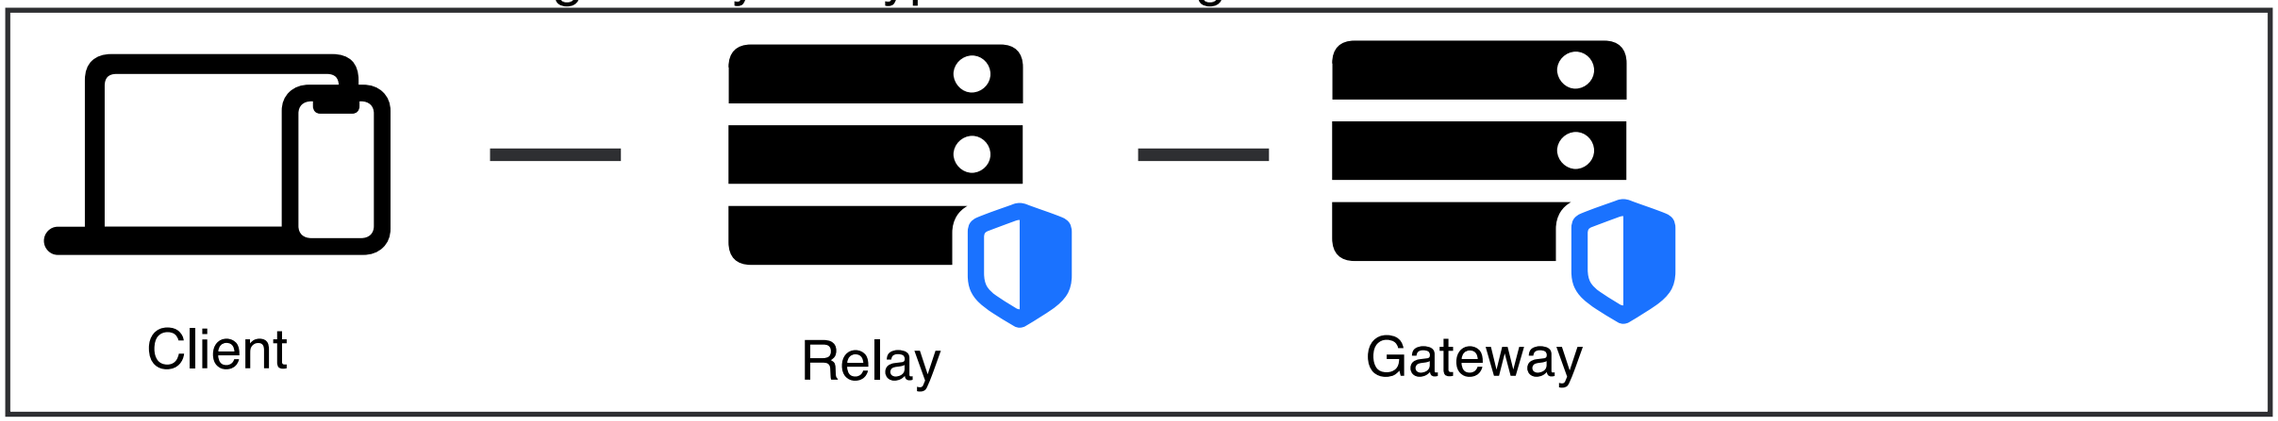

Context 3: Client-to-relay transport (+ client IP, etc)

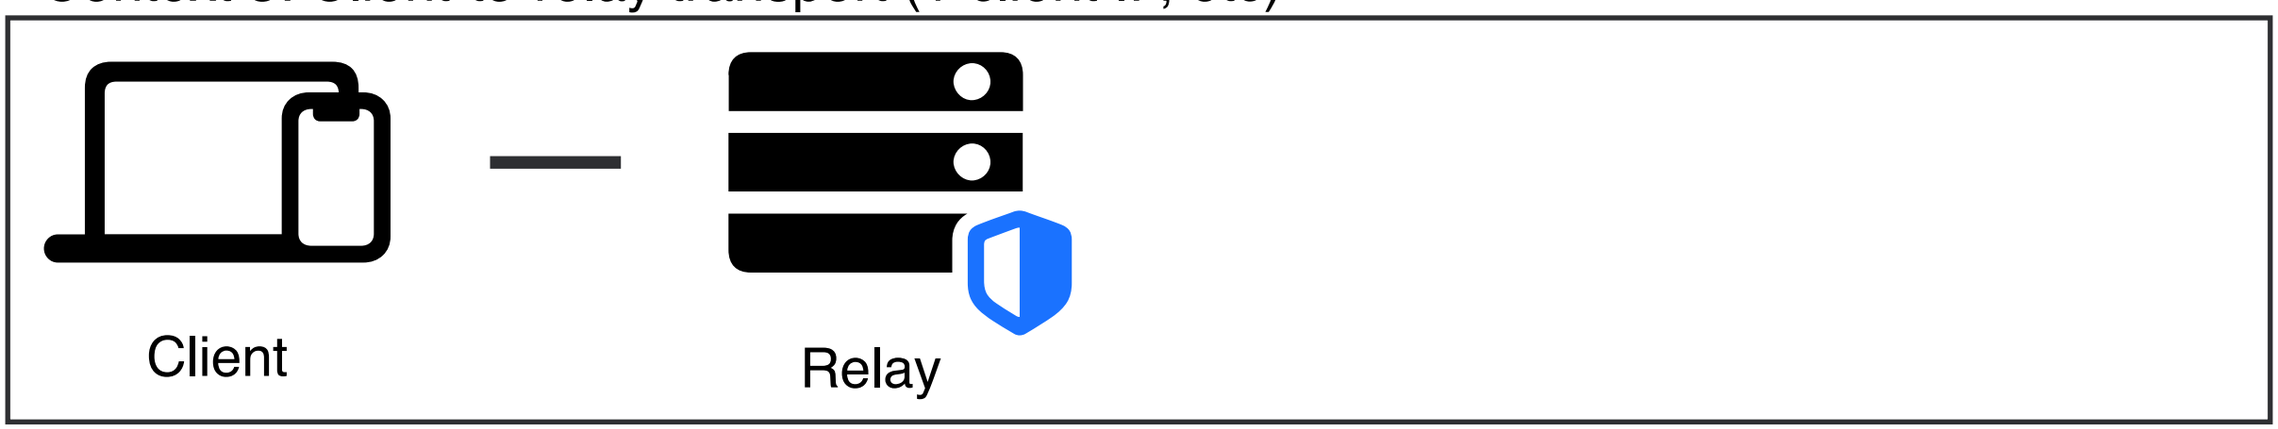

Observations

Partitioning is a **tool**, not a privacy panacea

- Relies on non-collusion across contexts
- Relies on careful selection of what data to include in a context
- Manageability needs to start relying on intentionally shared data
- Pay attention to performance in protocol design
- Partitioning doesn't solve traffic analysis without additional techniques being applied

Please read and provide input!