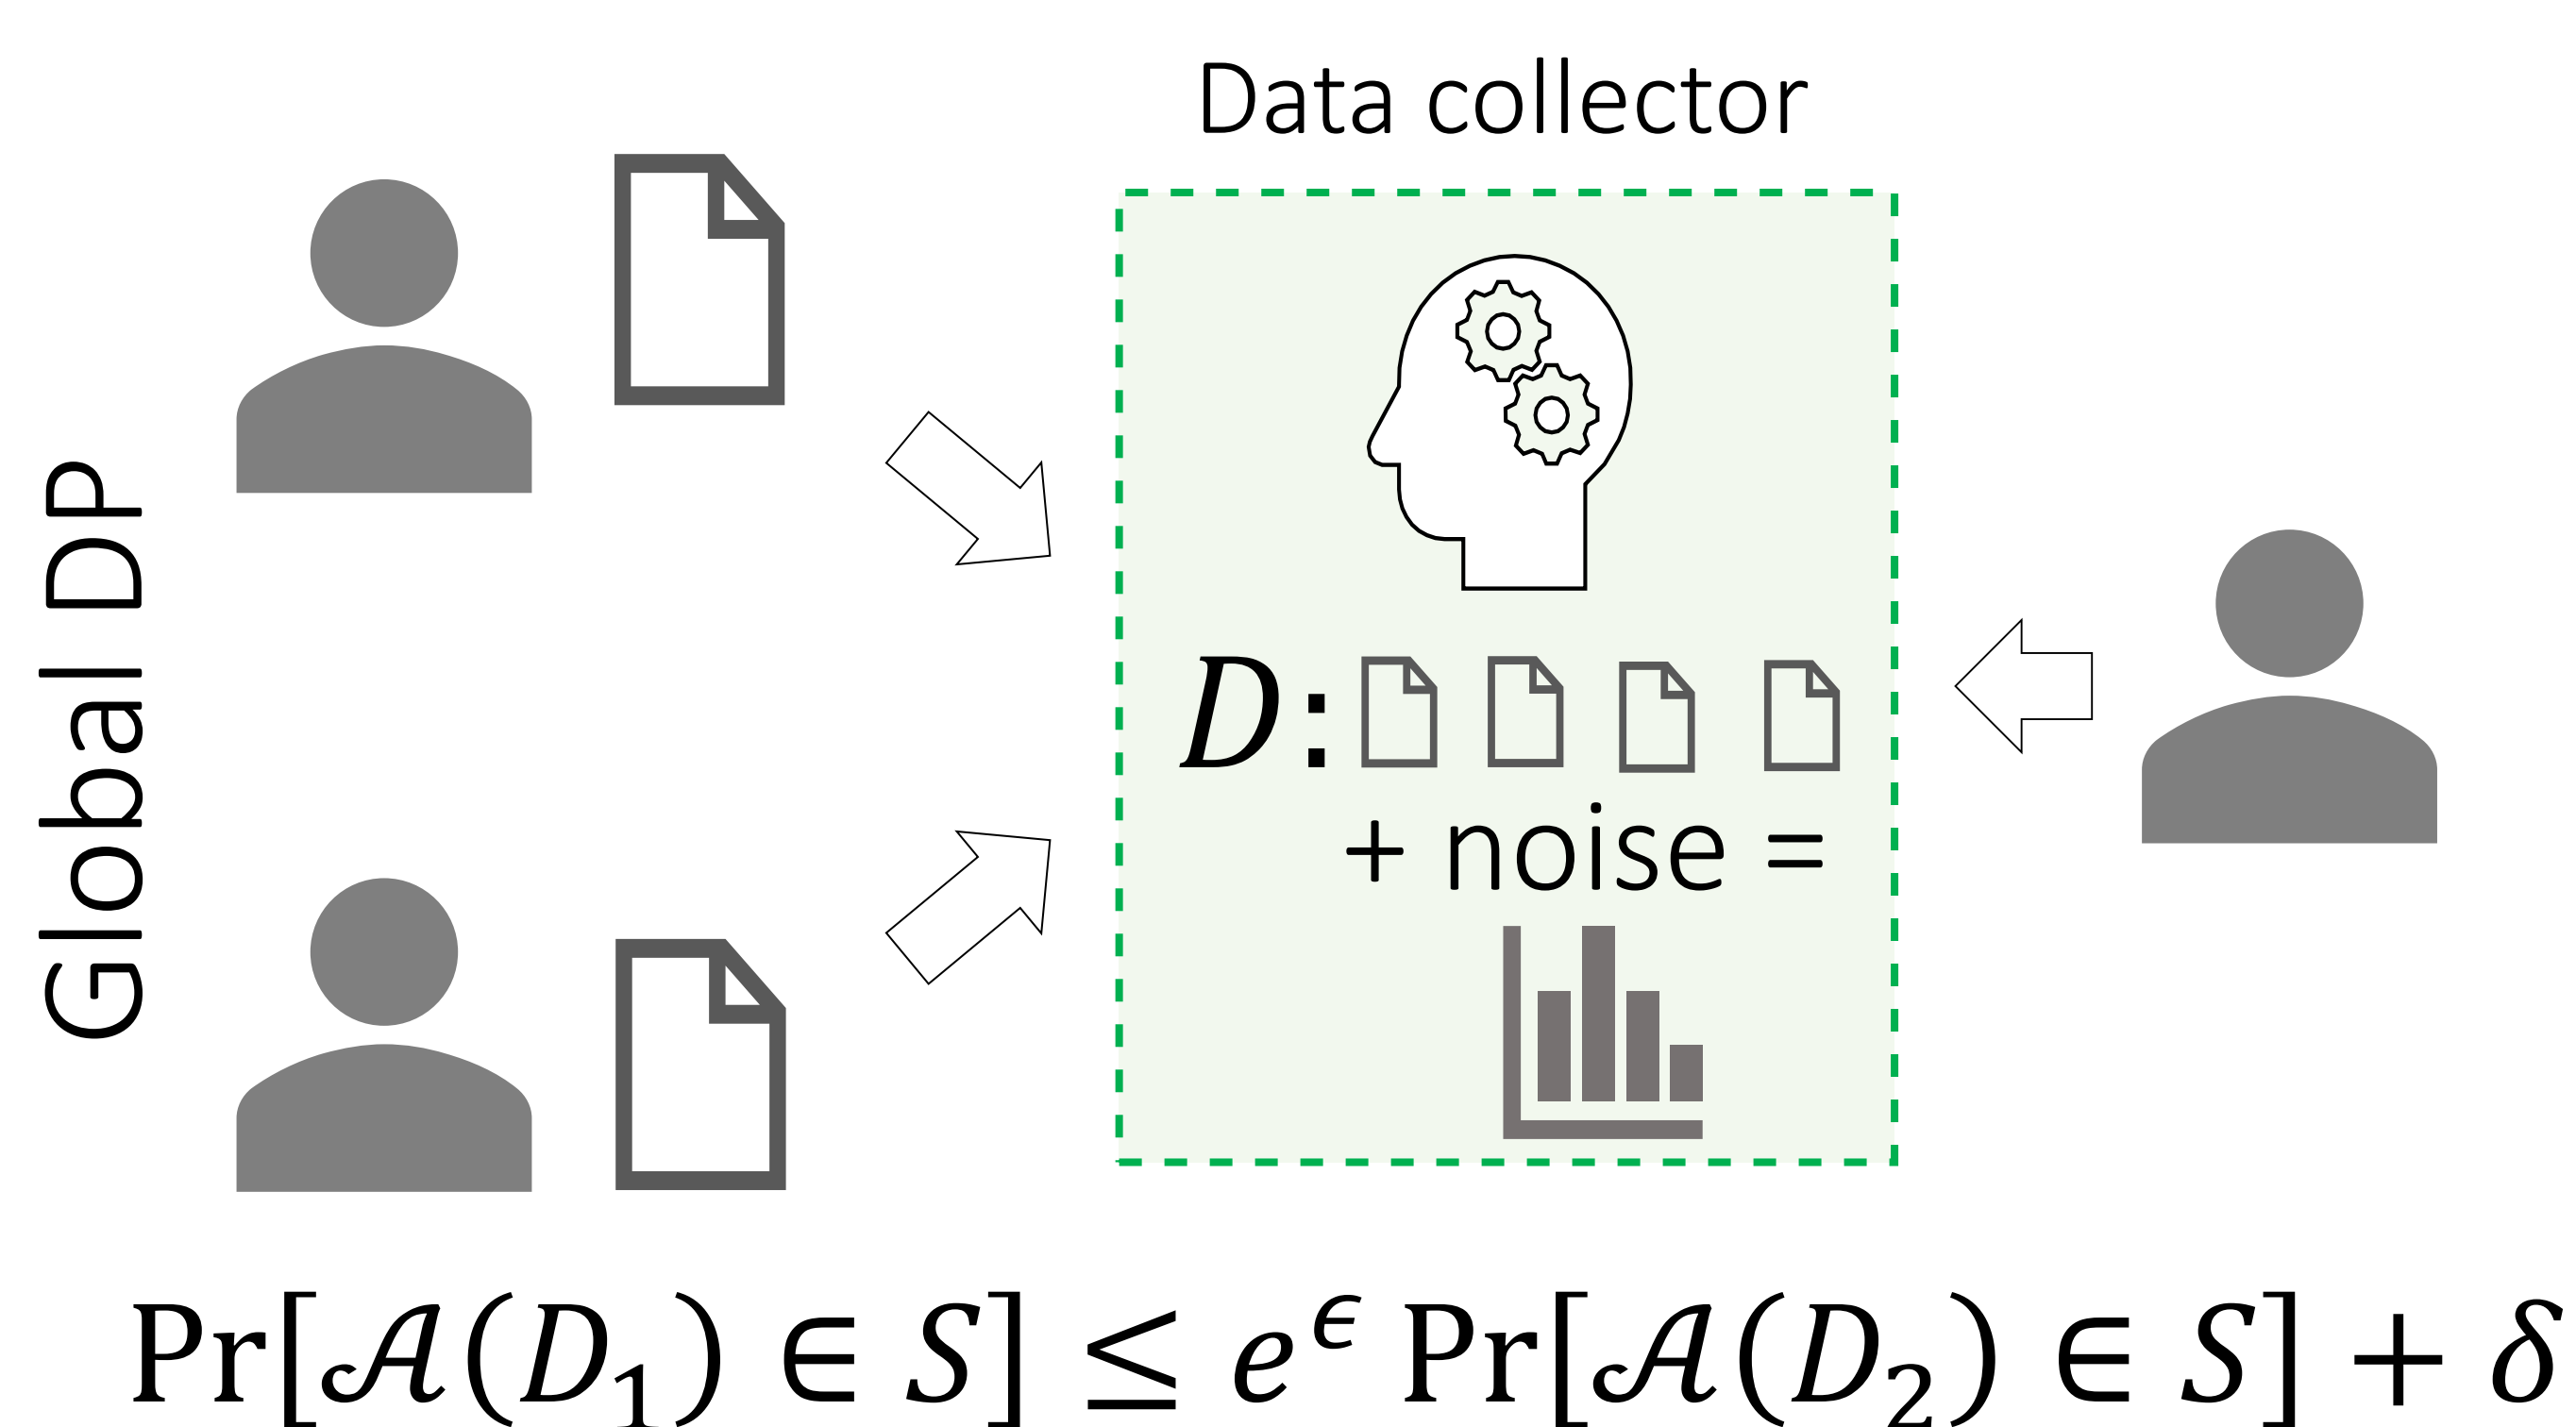
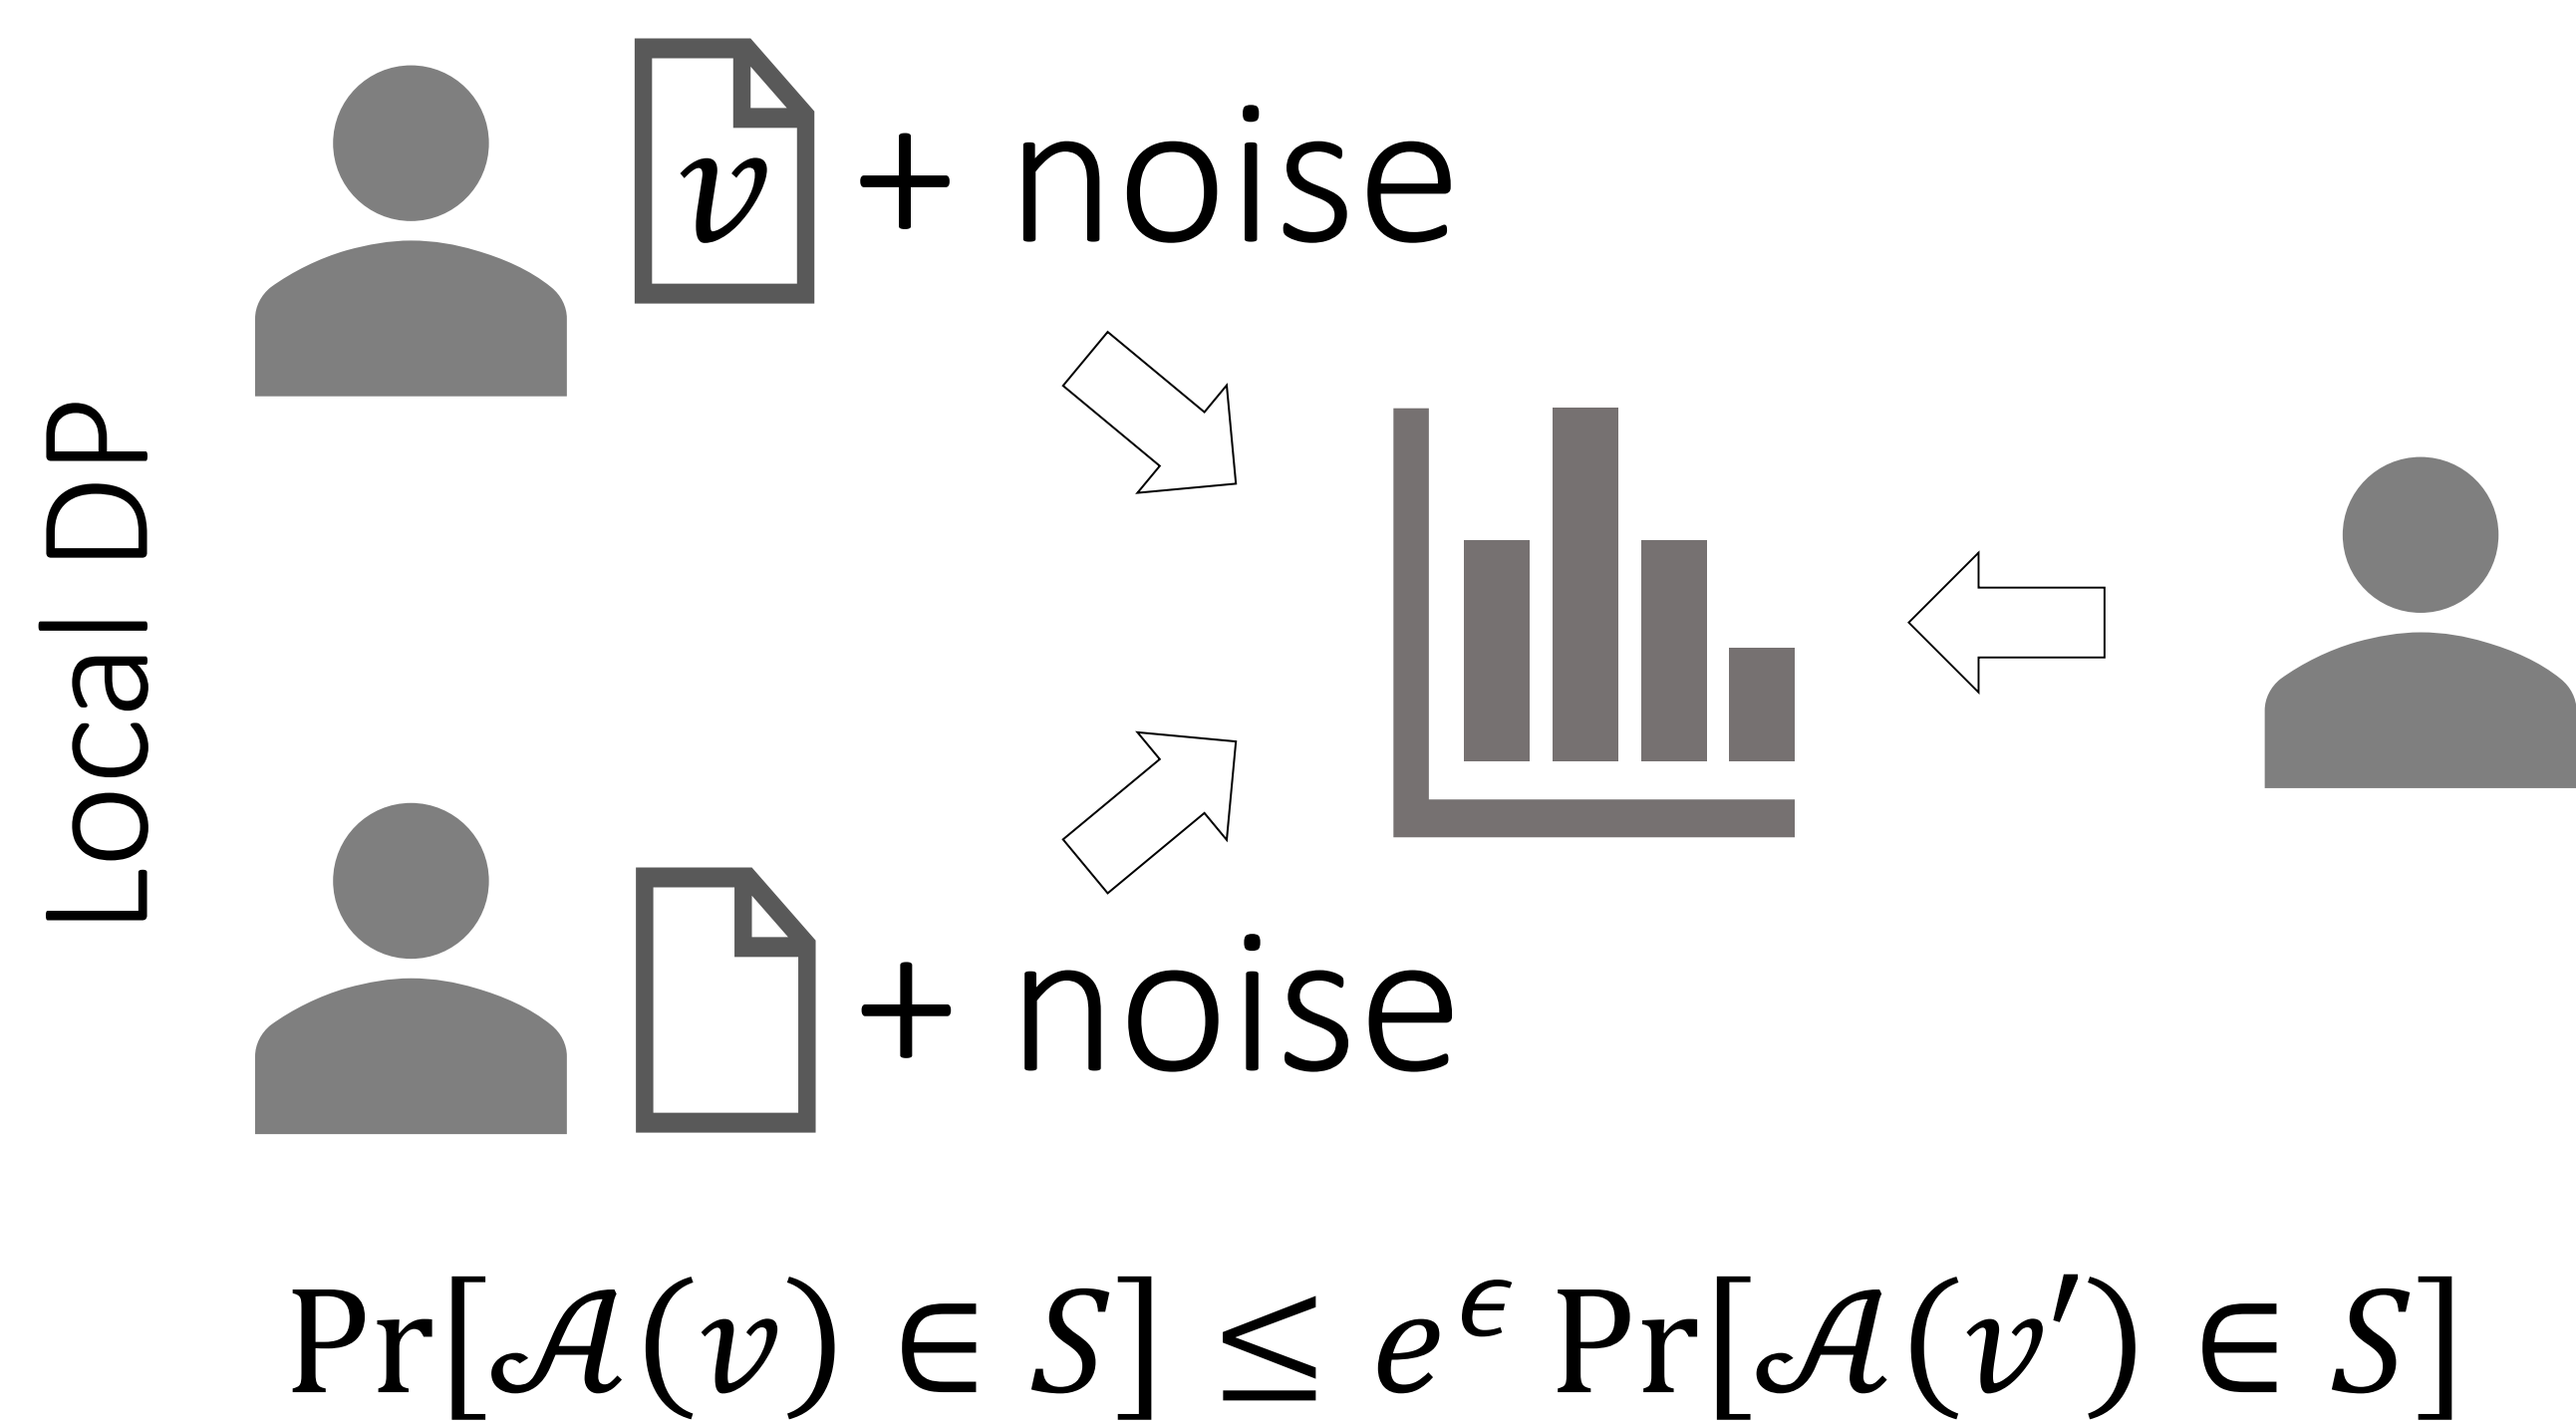


# An Algorithmic Framework For Differentially Private Data Analysis on Trusted Processors

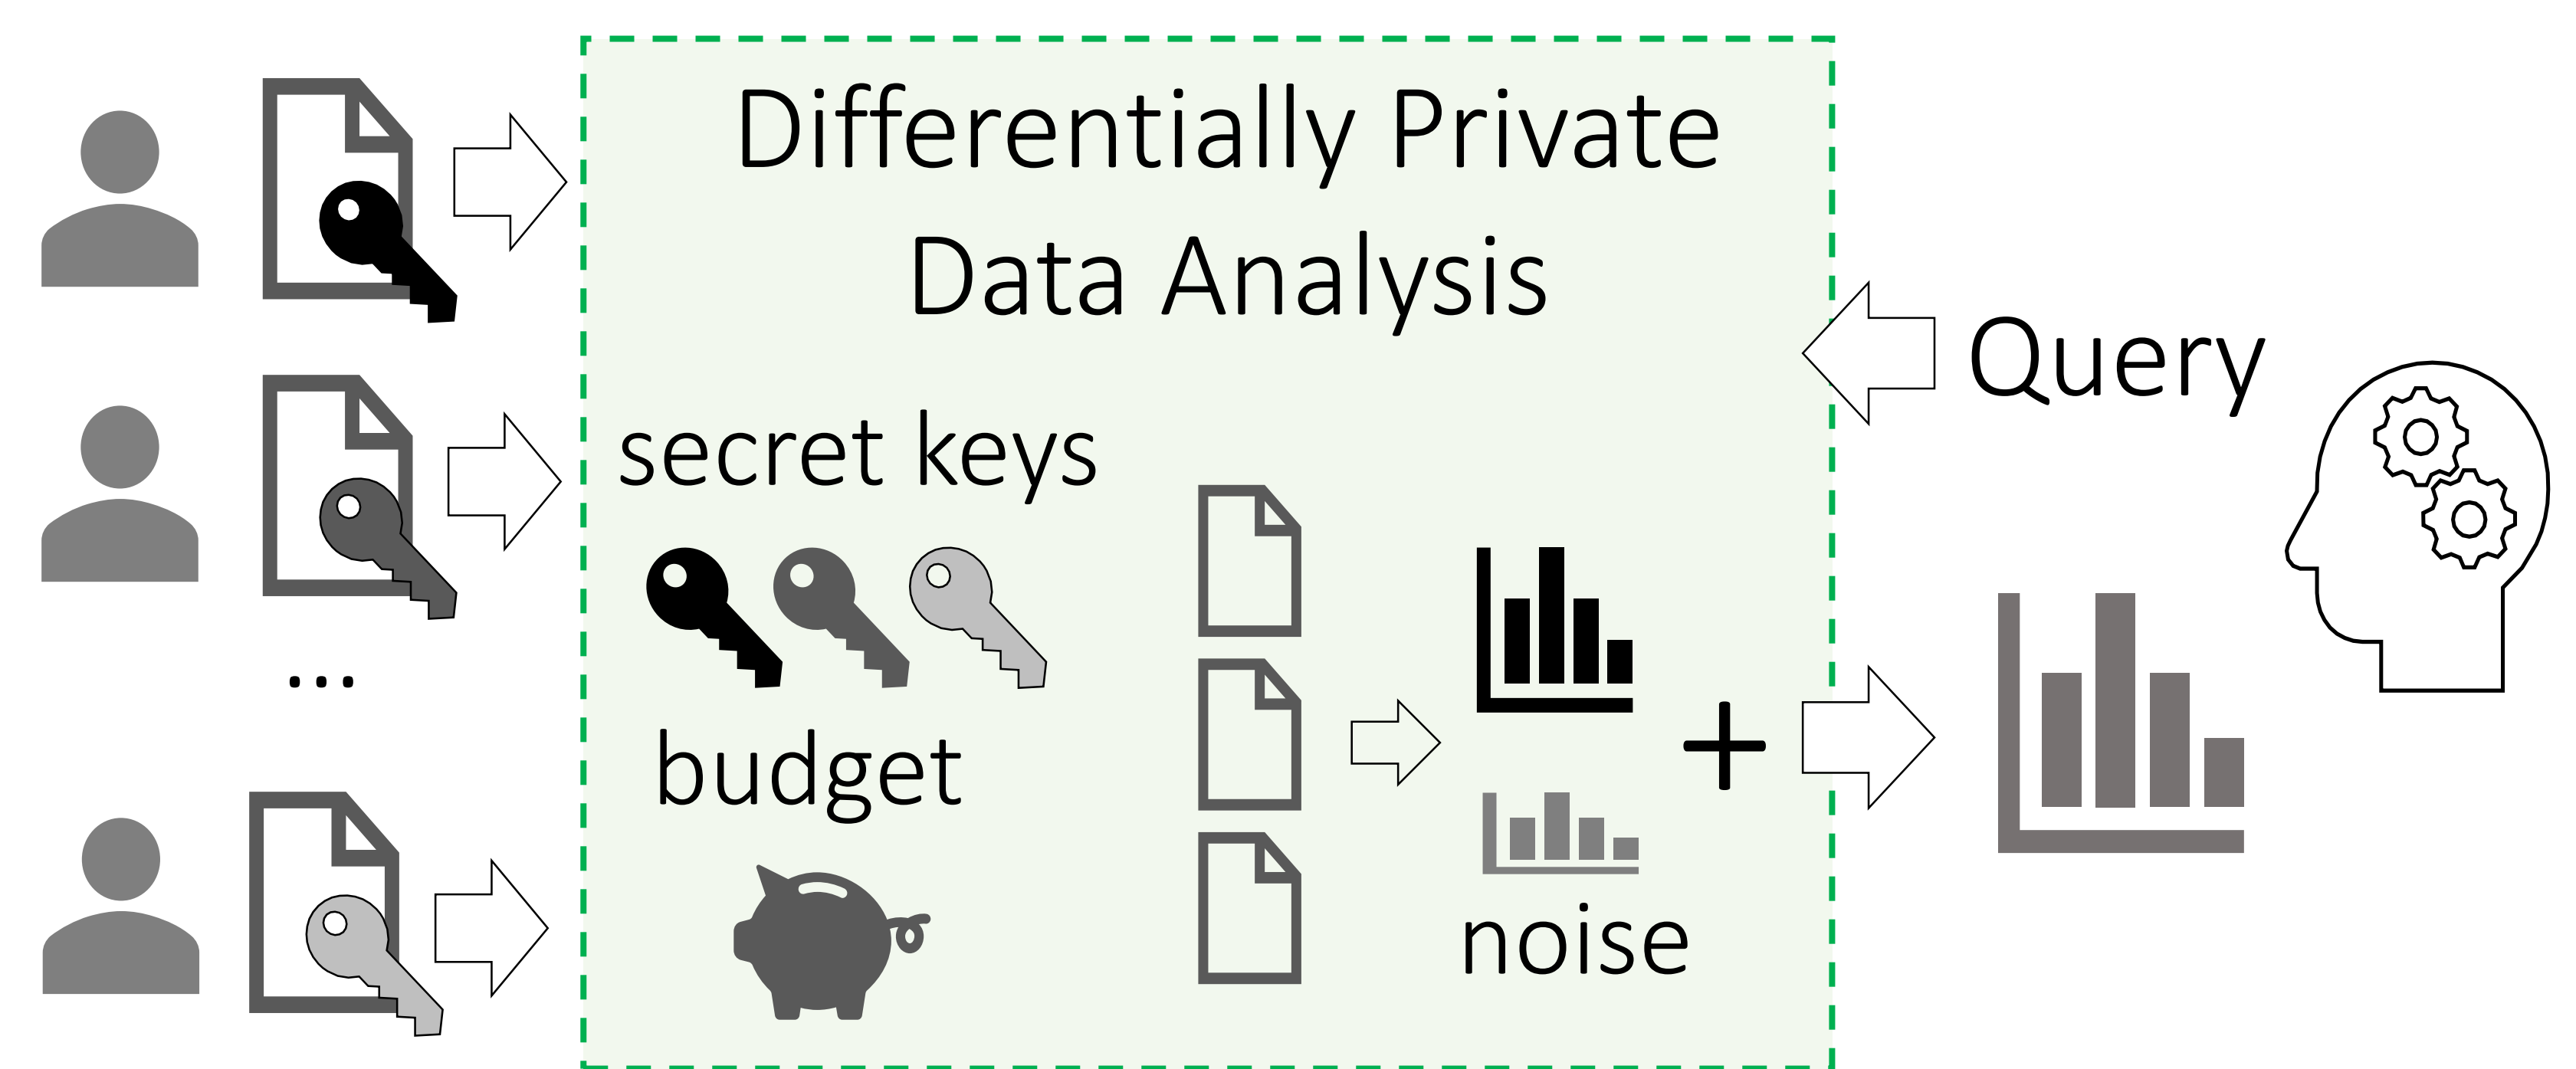
Joshua Allen, Bolin Ding, Janardhan Kulkarni

Harsha Nori, Olga Ohrimenko and Sergey Yekhanin

## Local vs. Global Differential Privacy (DP)



## Differential Privacy with Trusted Processors




### Trusted Execution Environment:

- containers for code and data
- isolated from the rest of the system (hypervisor, OS)
- data always encrypted in RAM
- remote attestation

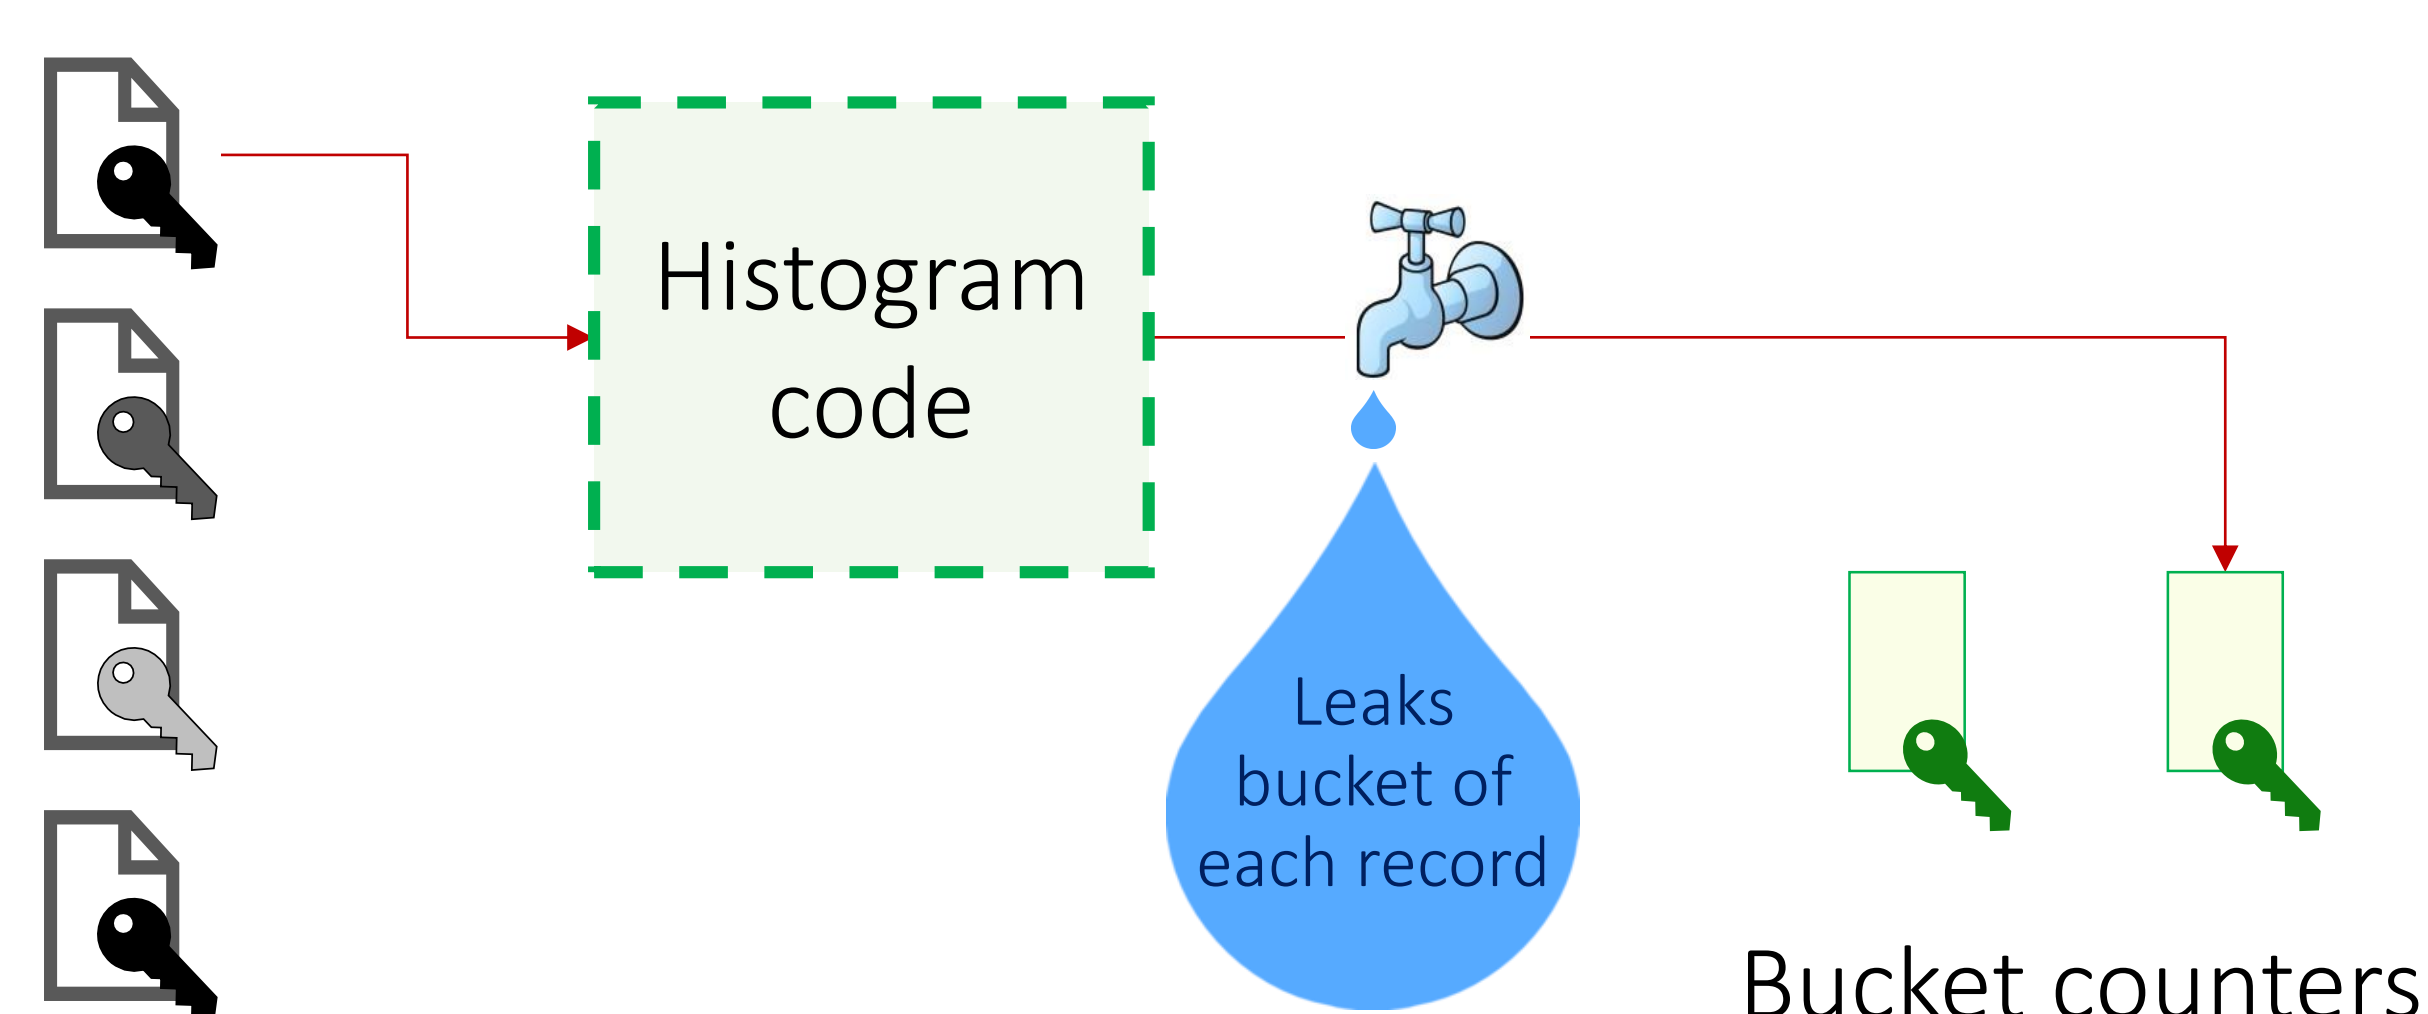


Intel SGX

## Information is Leaked via Side-Channels



Memory access patterns to external memory compromise differential privacy guarantees



## Oblivious Differential Privacy

$$\Pr[\mathcal{A}(D_1) \in (O, S)] \leq e^\epsilon \Pr[\mathcal{A}(D_2) \in (O, S)] + \delta$$

where  $O$  is a subset of outputs and  $S$  is a subset of memory access patterns produced by  $\mathcal{A}$

### Oblivious Differentially Private Histogram Algorithm:

