IronSafe

A Secure and Policy-Compliant Query Processing Architecture

Harshavardhan Unnibhavi

David Cerdeira, Antonio Barbalace, Nuno Santos, Pramod Bhatotia

THE UNIVERSITY of EDINBURGH

TÉCNICO LISBOA

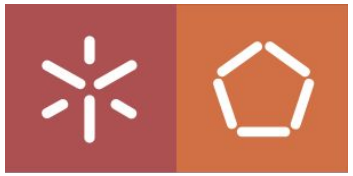

Universidade do Minho
Escola de Engenharia

ACM SIGMOD 2022

Data analytics in the cloud

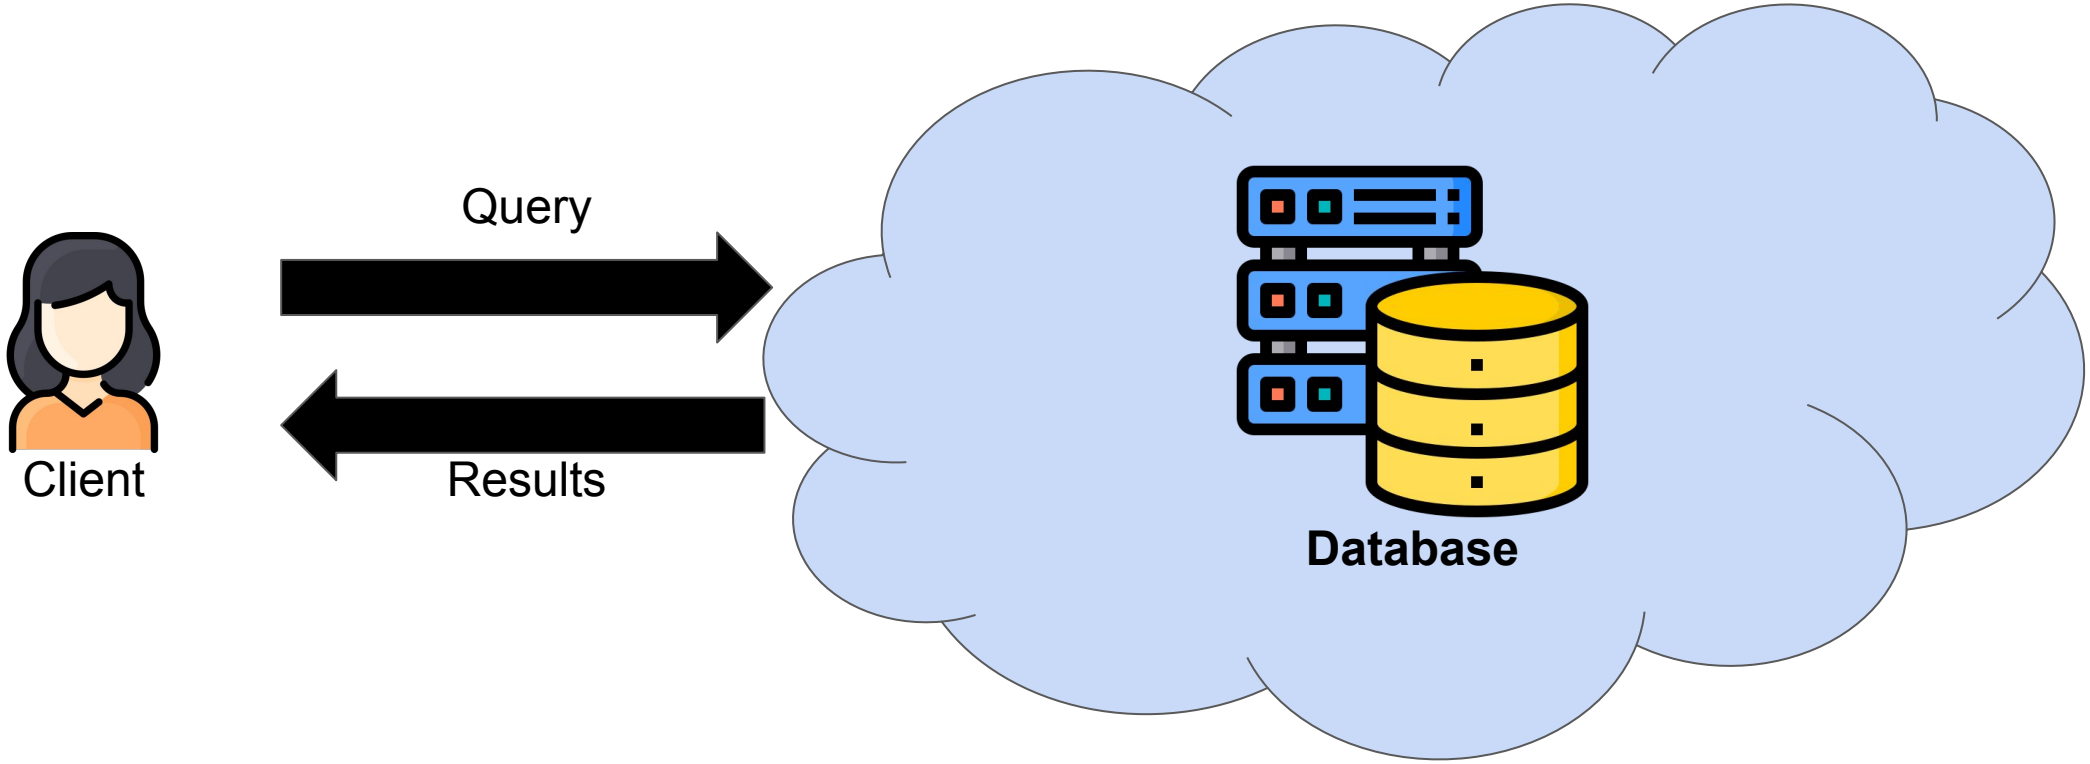

Cloud enables scalable and fault-tolerant computing in a cost-effective manner

The challenges of cloud computing

#1: Security

#2: Policy compliance

#1: Security challenges

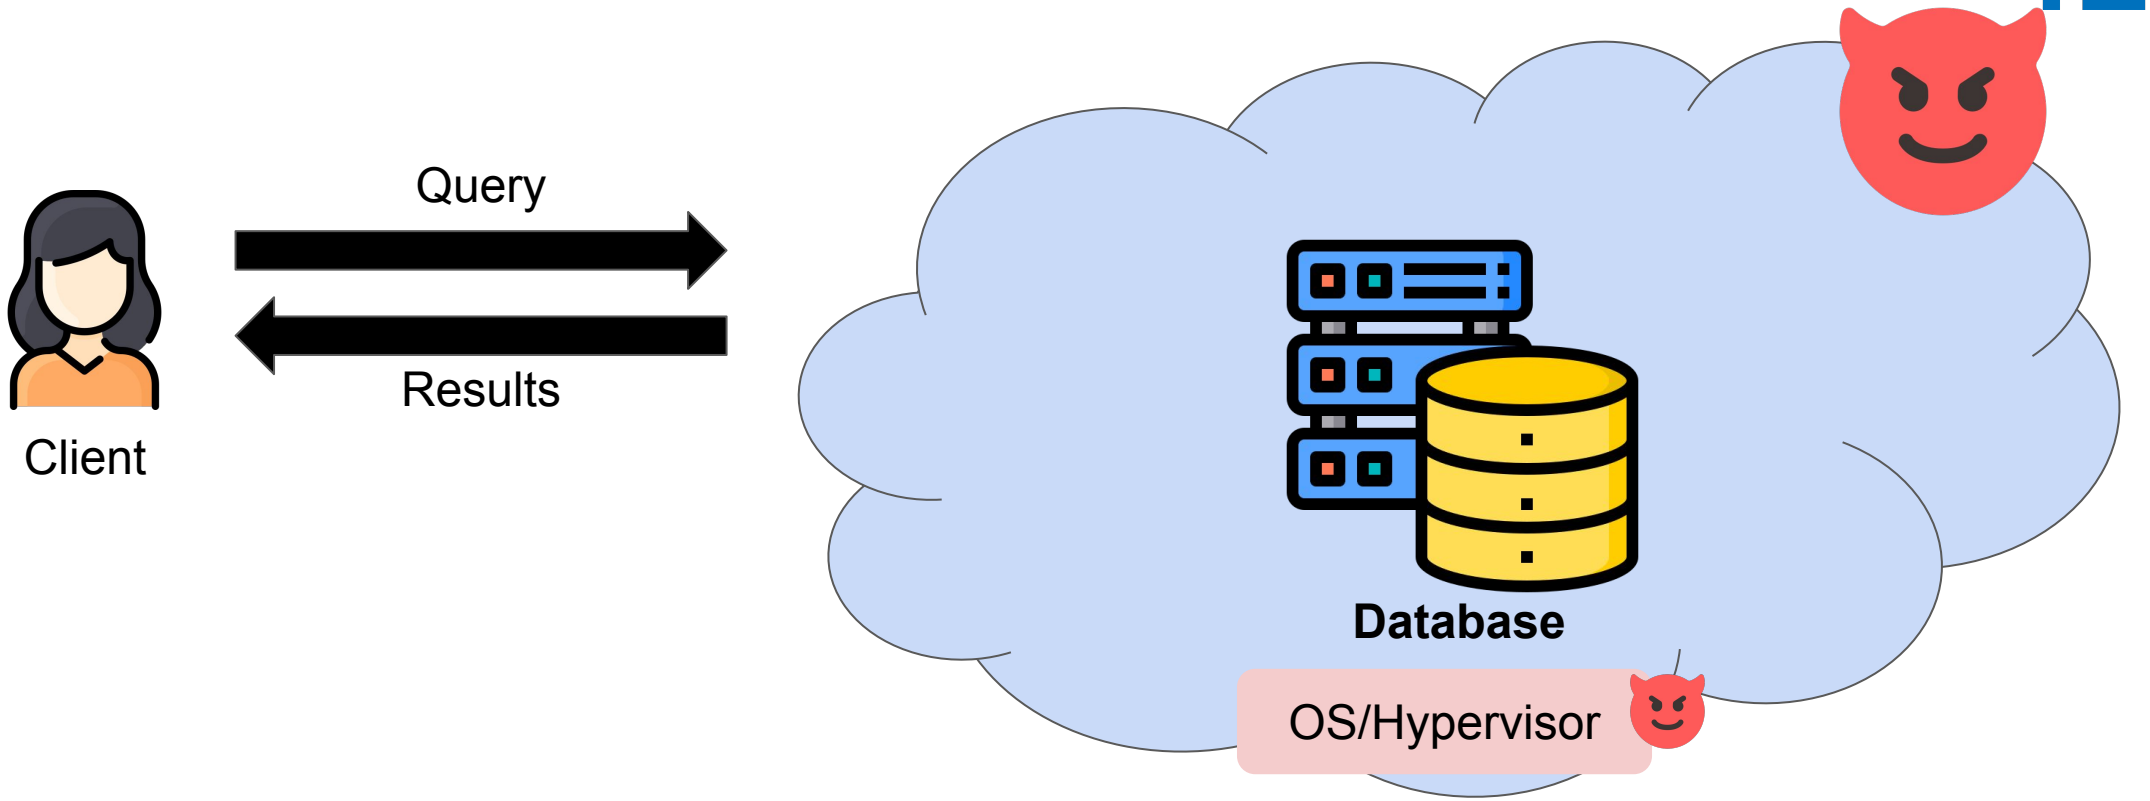

How do we ensure security in untrusted cloud environments?

#2: Policy compliance challenges

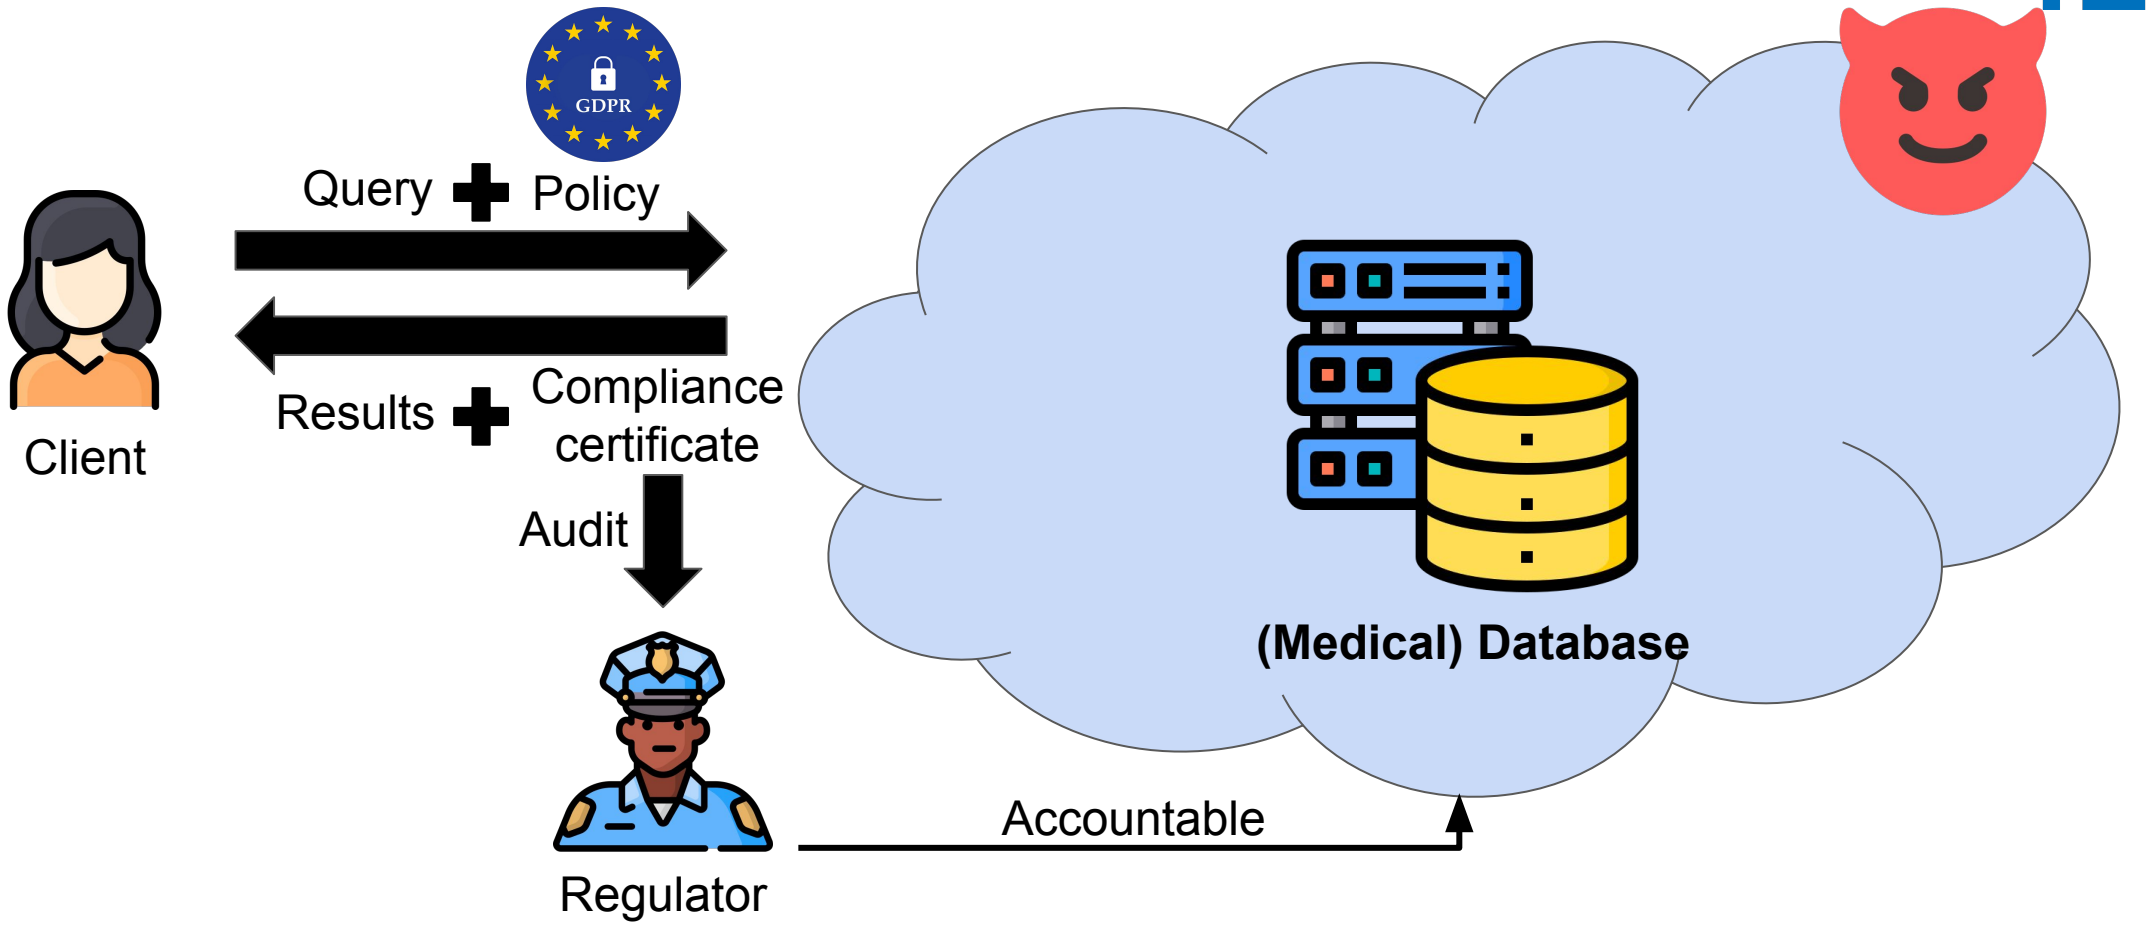

How do we ensure policy-compliant query processing that is auditable by a regulator?

Problem statement

To design a **secure, policy-compliant and high-performant** query processing architecture for untrusted cloud environments

Our proposal

IronSafe
A secure and policy-compliant query processing architecture

Security

Policy compliance

High performance

Key building blocks

IronSafe

Hardware-assisted trusted computing

+

Policy language and compliance infrastructure

+

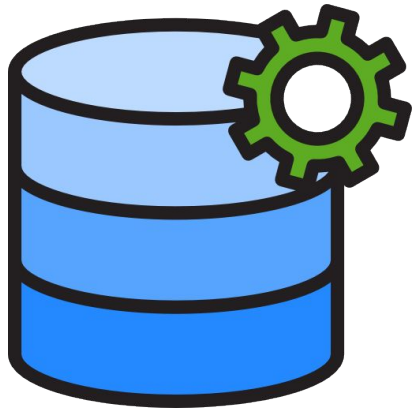

Near data processing (NDP)

IronSafe overview

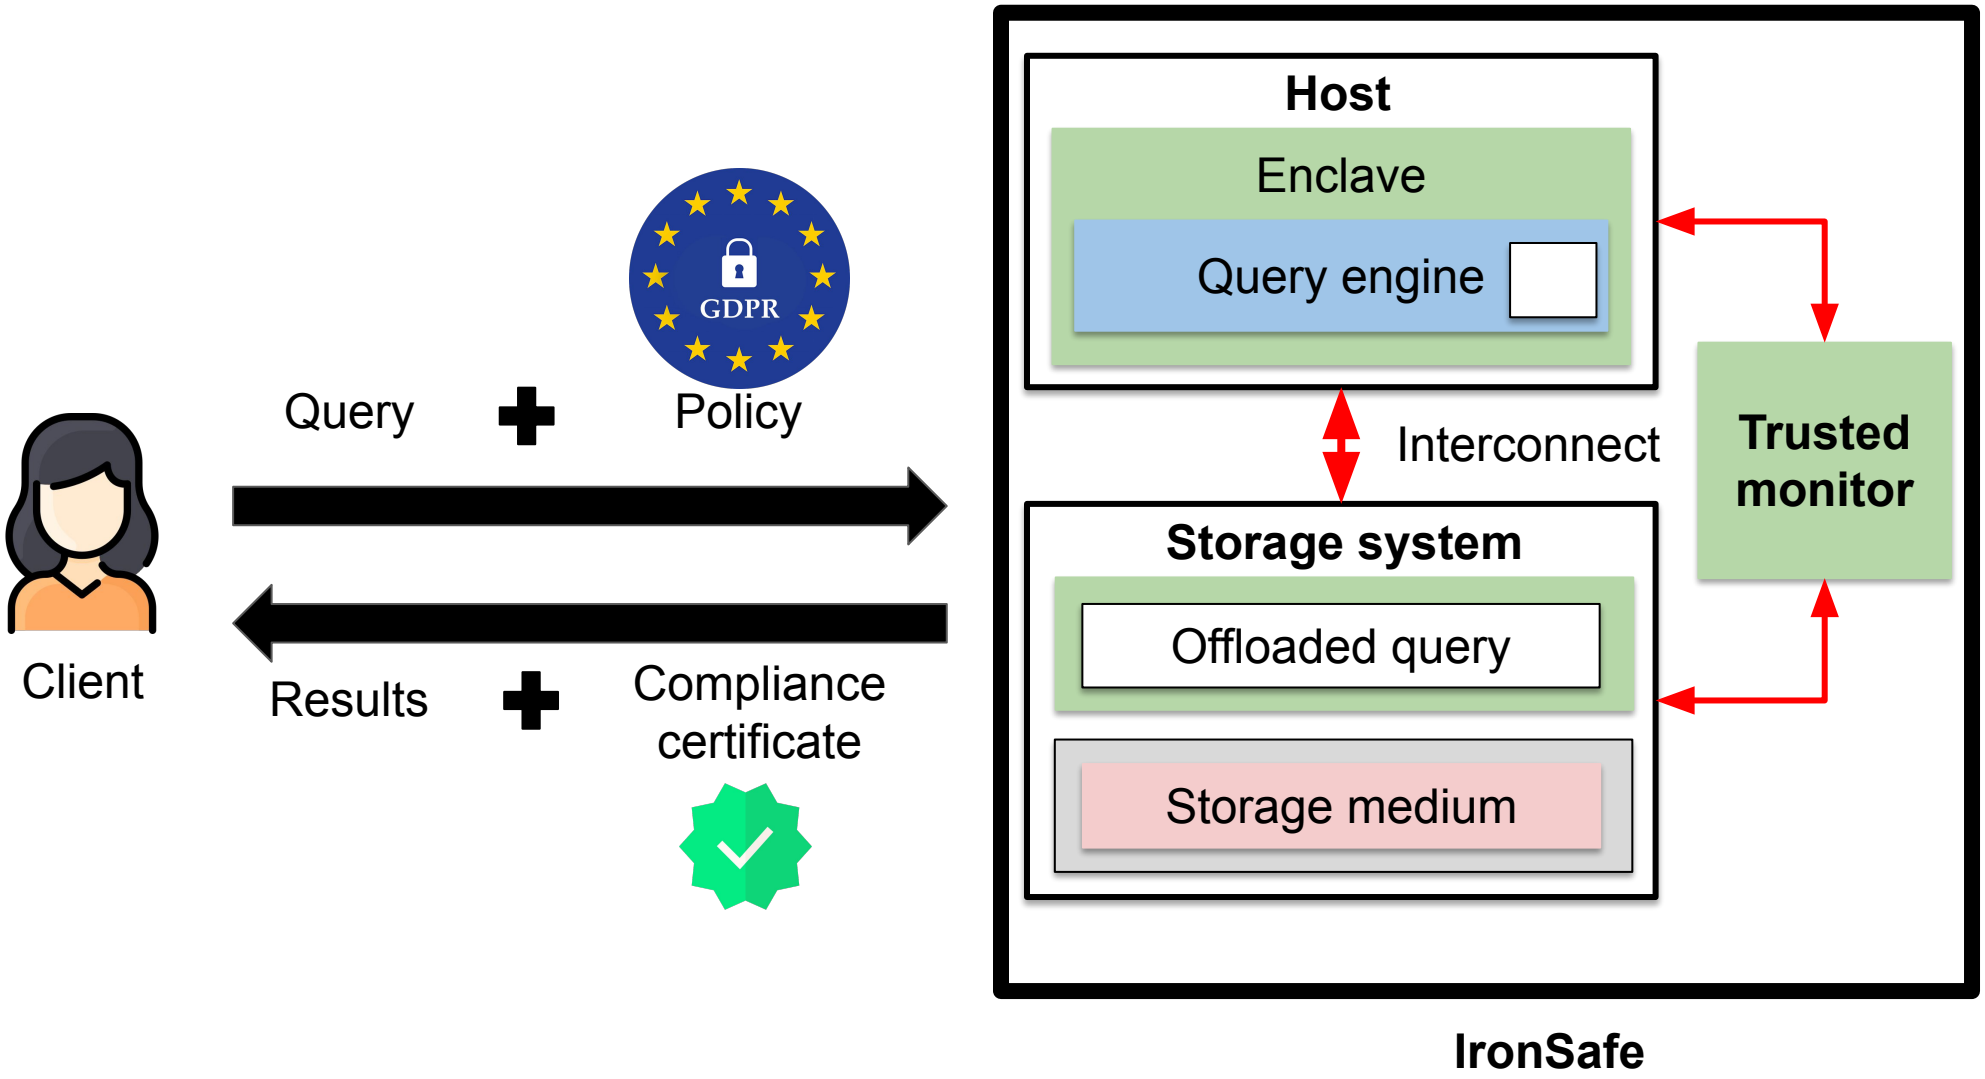

Outline

- ~~Motivation~~
- Background
- Design challenges
- Workflow
- Evaluation

Trusted computing

- Abstraction
  - Trusted Execution Environment (TEE)
  - Rich Execution Environment (REE)

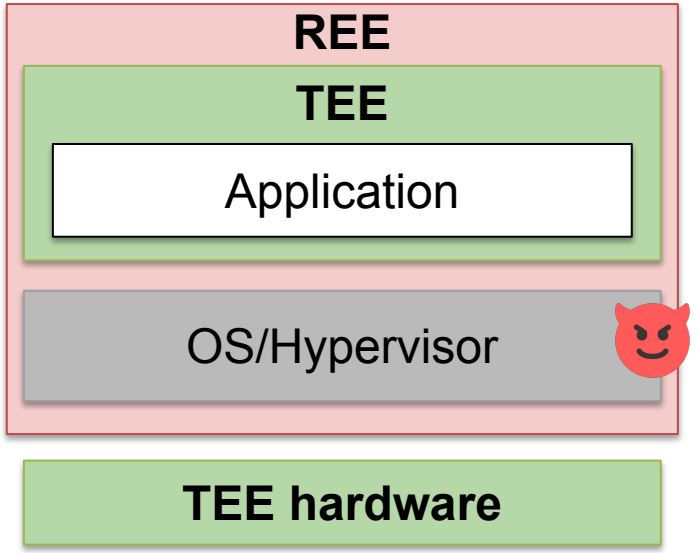

- Remote attestation
  - Authenticate hardware and software

IronSafe leverages trusted computing to provide strong security guarantees

Near data processing (NDP)

- Heterogeneous deployment
  - Specialized accelerators for data processing
  - x86 host and ARM storage system

- Near data processing
  - Offload computation to storage

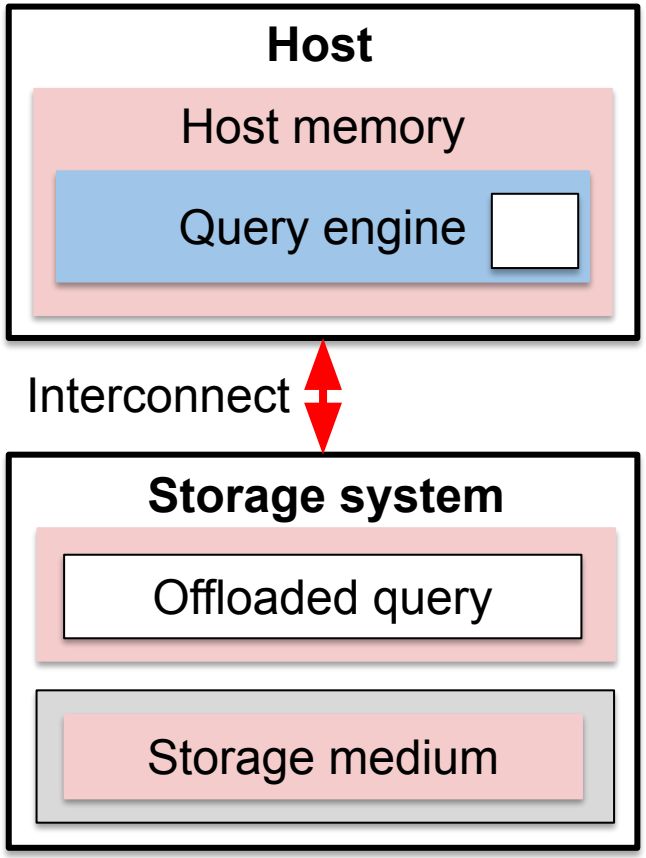

IronSafe leverages NDP to achieve high performance

Outline

- ~~Motivation~~
- ~~Background~~
- Design challenges
- Workflow
- Evaluation

Challenge #1: Heterogenous TEEs

- Heterogeneous host and storage
  - Heterogeneous ISAs and TEEs
  - Different security guarantees

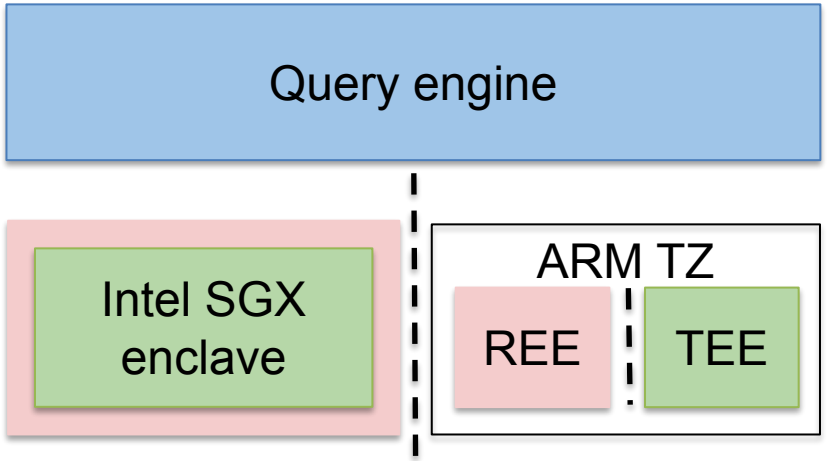

- Combine execution across TEEs
  - Host: Intel SGX
  - Storage: ARM TrustZone

A heterogeneous confidential computing framework

#1: A heterogeneous confidential computing framework

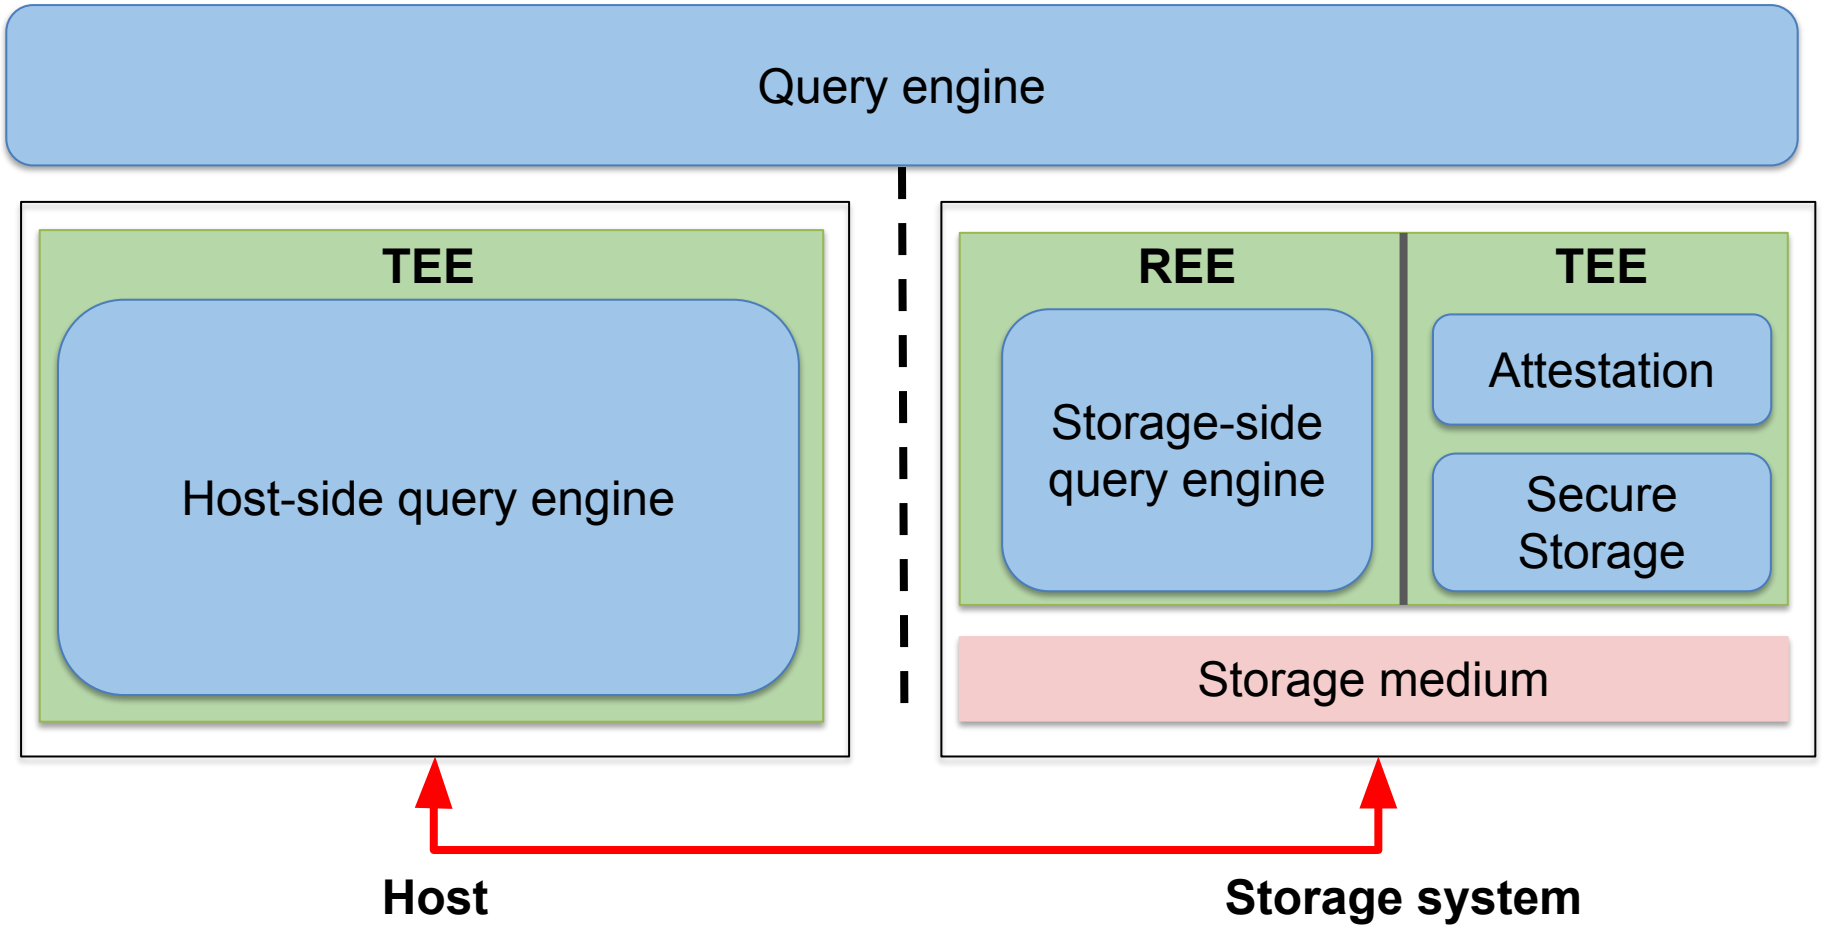

Challenge #2: Policy specification and compliance

- Policy specification
  - Expressible but simple and less error-prone

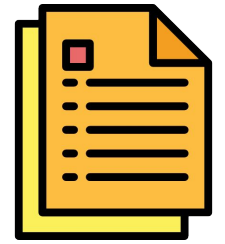

- Policy enforcement
  - Guarantee integrity of all components for every query
  - Multiple versions of hardware and software

- Proof of policy compliance
  - Guarantee compliance for every query

A policy compliance monitor

#2: Policy compliance monitor

- Policy specification
  - IronSafe's declarative policy language
- Policy enforcement
  - Attestation of both host and storage system
- Proofs of policy compliance
  - Cryptographic certificates signed by trusted monitor

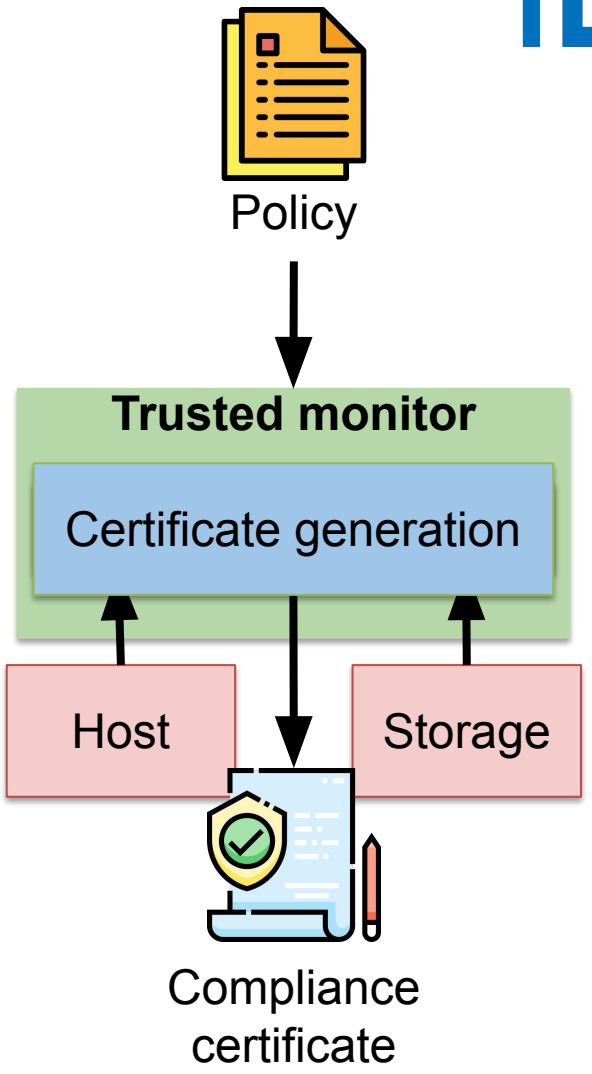

Outline

- ~~Motivation~~
- ~~Background~~
- ~~Design challenges~~
- Workflow
- Evaluation

IronSafe workflow

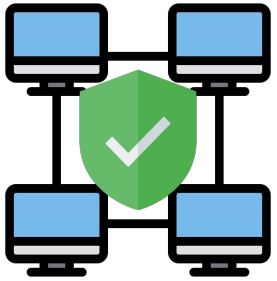

#1

Trust establishment

#2

Policy compliance

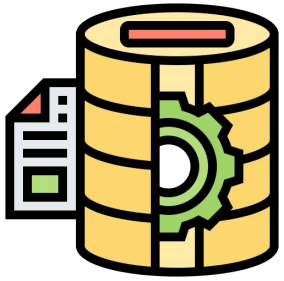

#3

Query execution

#1: Trust establishment

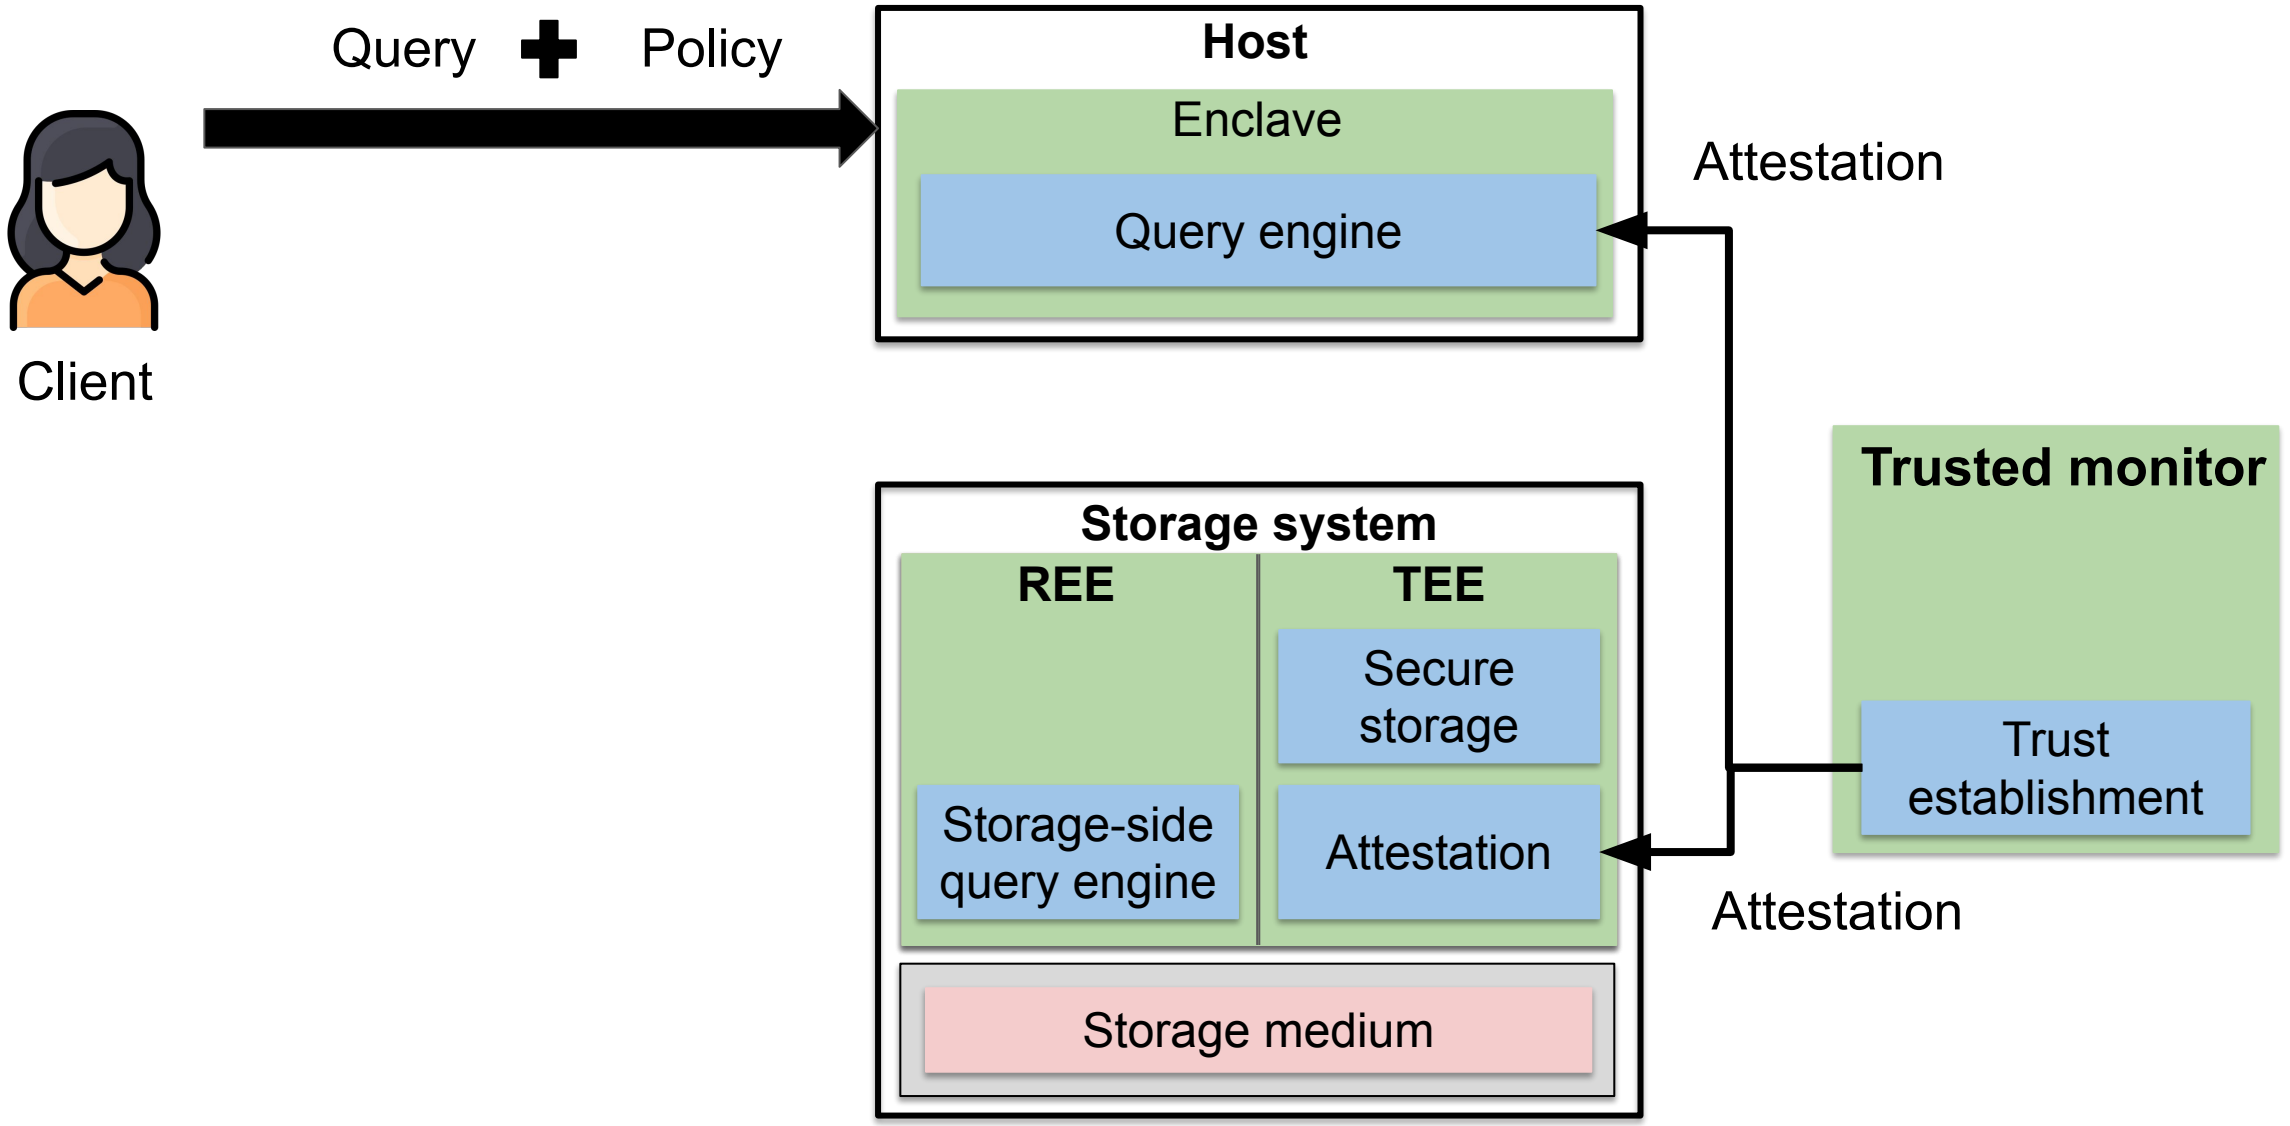

#2: Policy enforcement

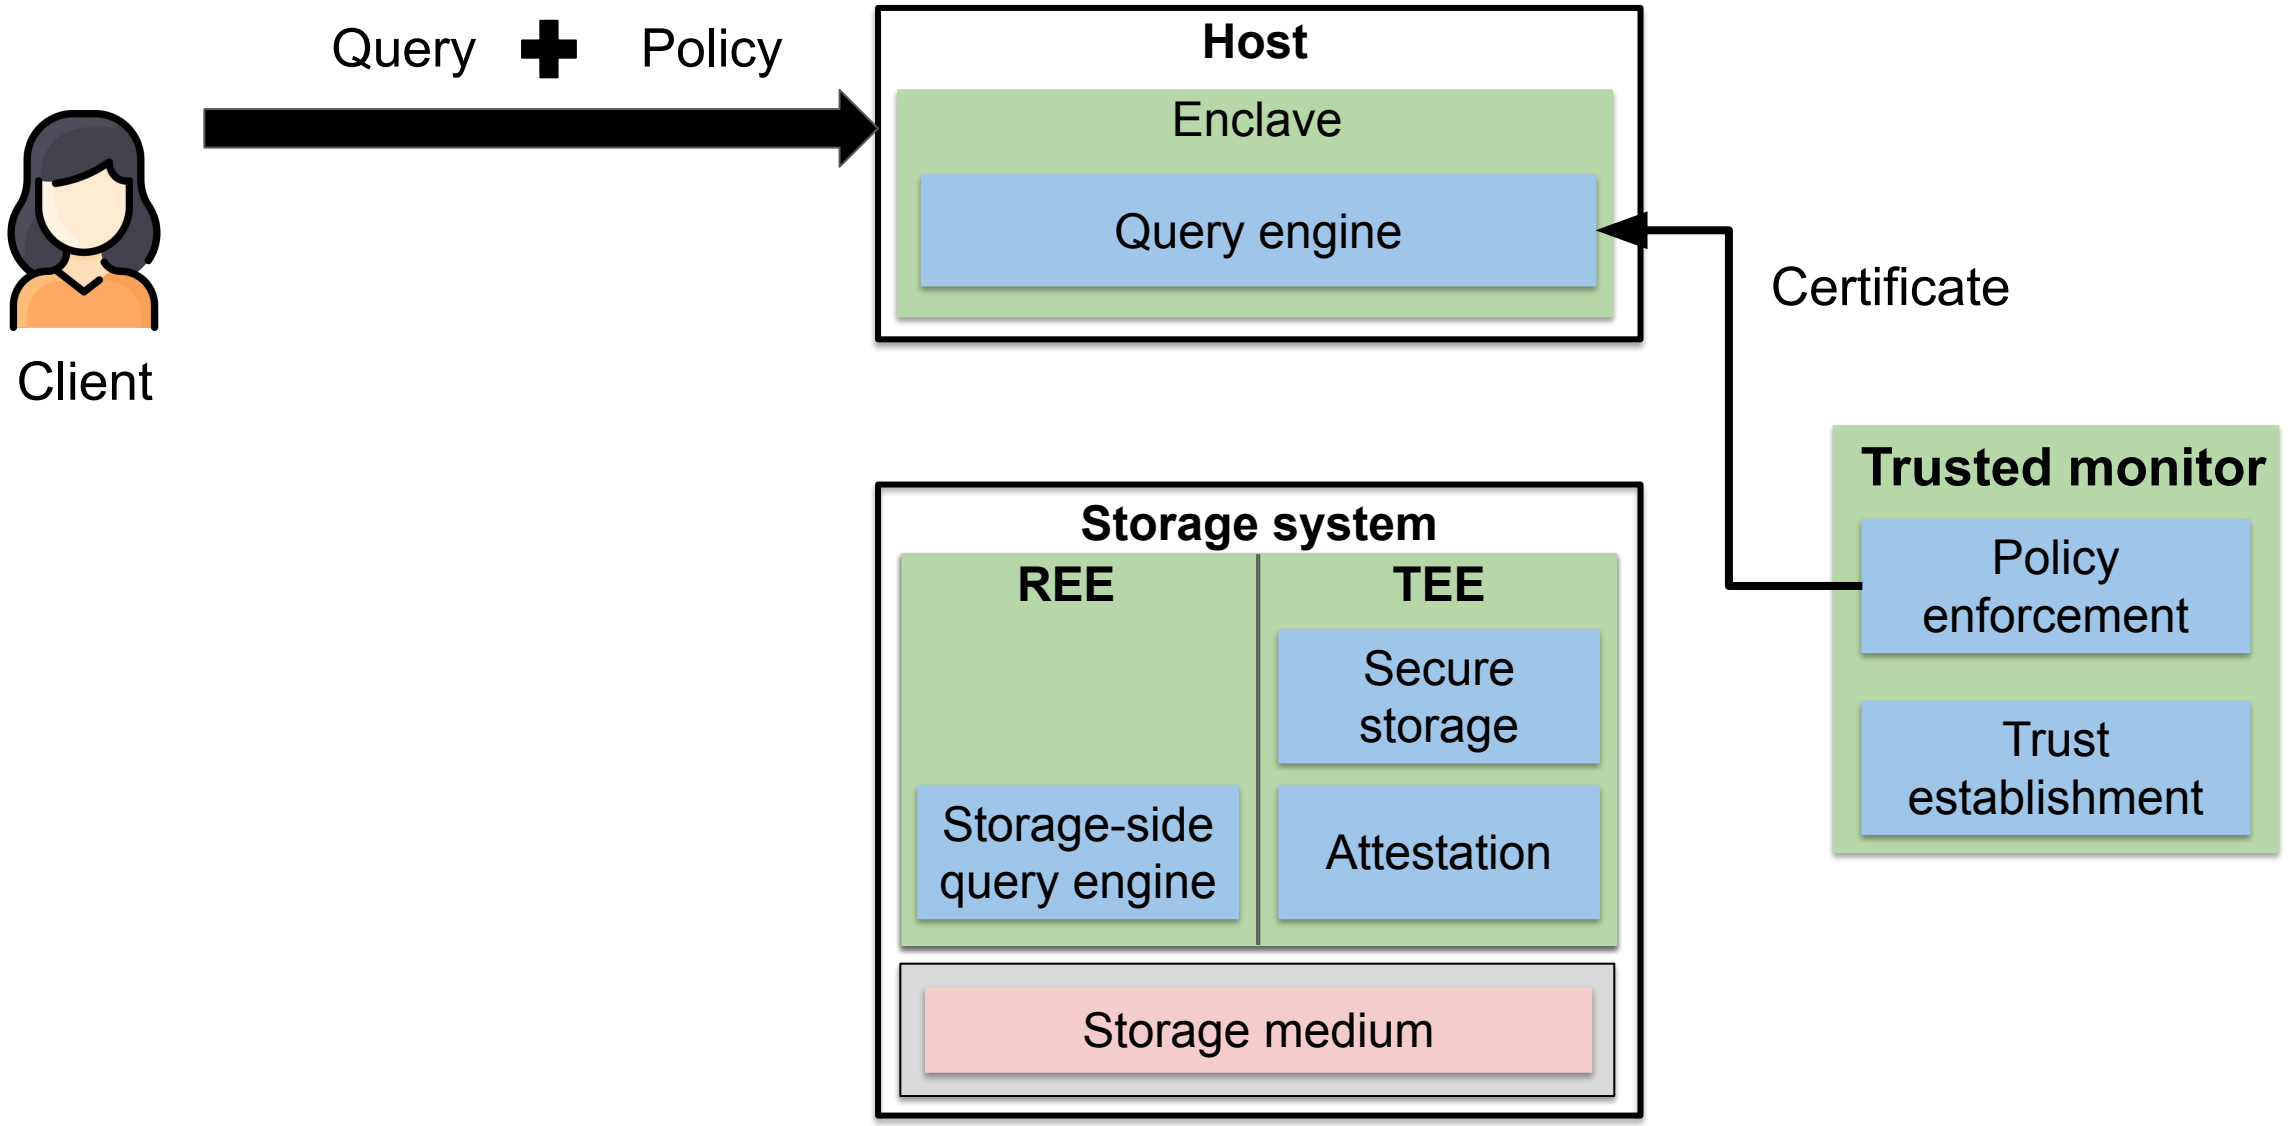

#3: Query execution

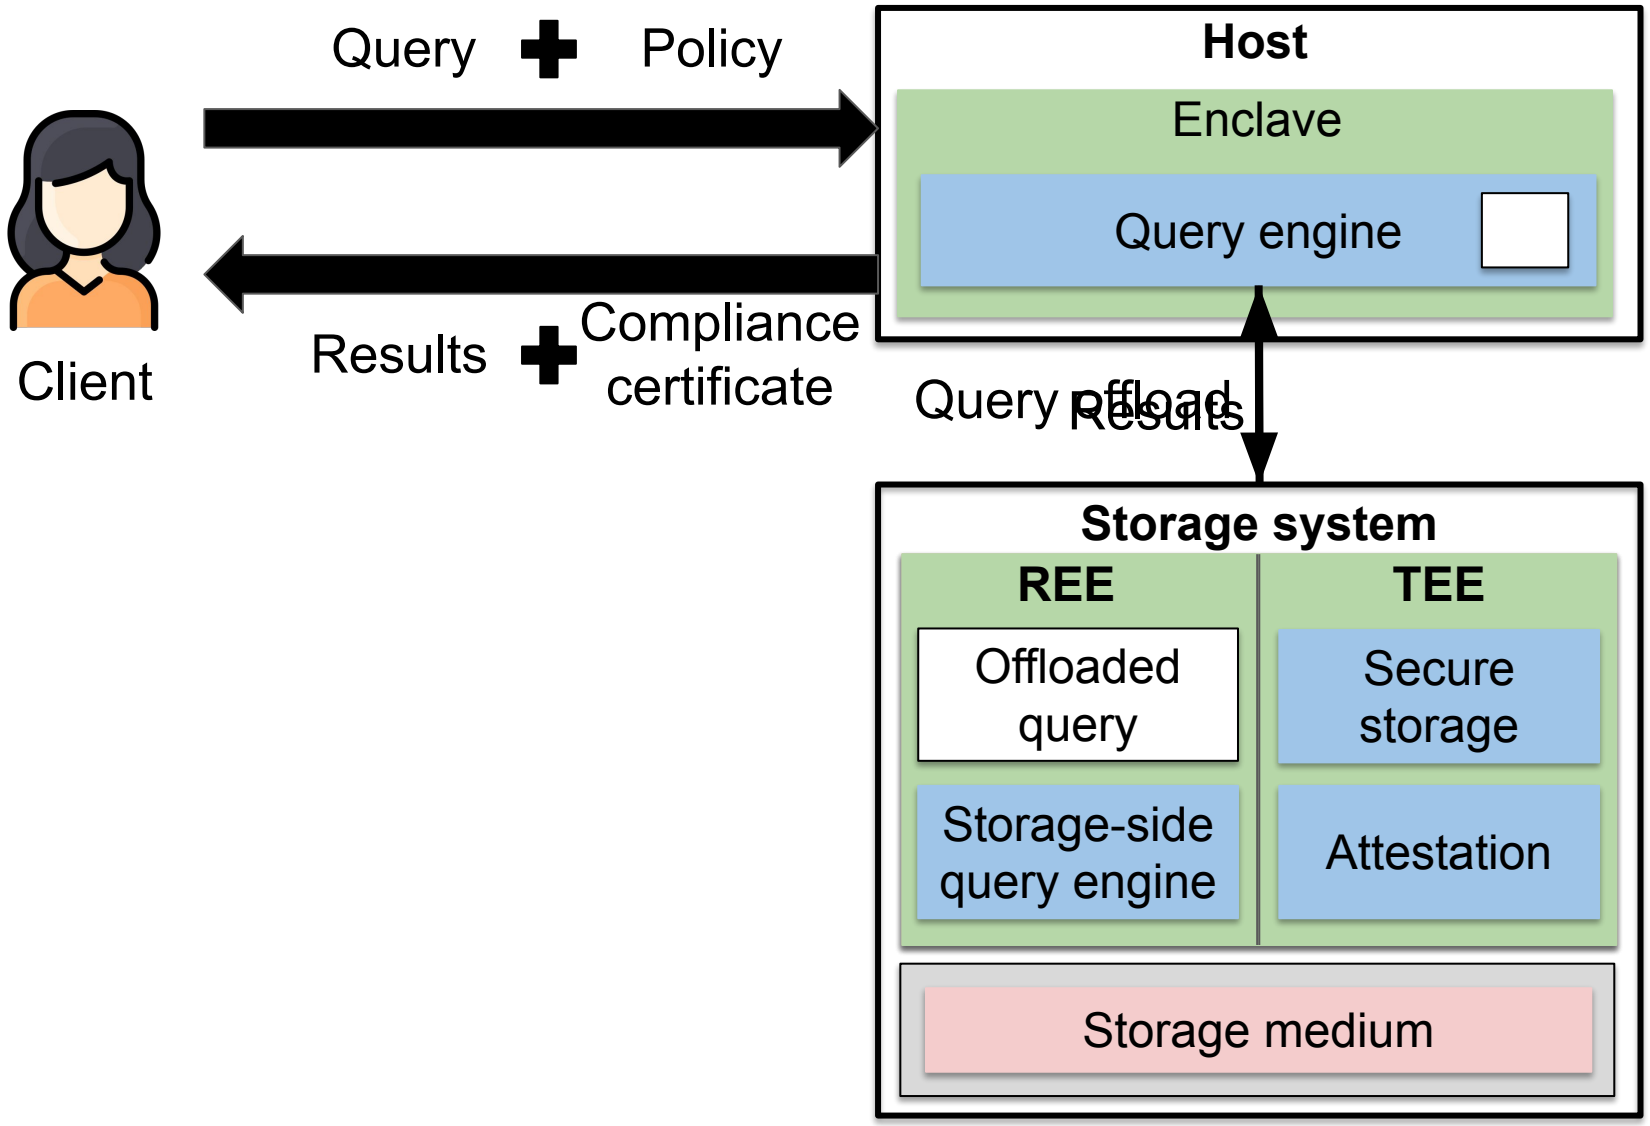

Trusted monitor

- Policy enforcement
- Trust establishment

Outline

- ~~Motivation~~
- ~~Background~~
- ~~Design challenges~~
- ~~Workflow~~
- Evaluation

Evaluation

- Questions
  1) Performance gains
  2) Policy compliance use cases
- Workloads
  - SQL queries: TPC-H benchmark
  - GDPR anti-patterns (policy compliance)
- Setup
  - Host: i9-10900K, 64GiB memory
  - Interconnect: 40GbE
  - Storage: 16 core ARM cortex A72, 32 GiB memory, 1TB M.2 NVMe

Q1: Performance gains

IronSafe speedup compared to secure host only processing

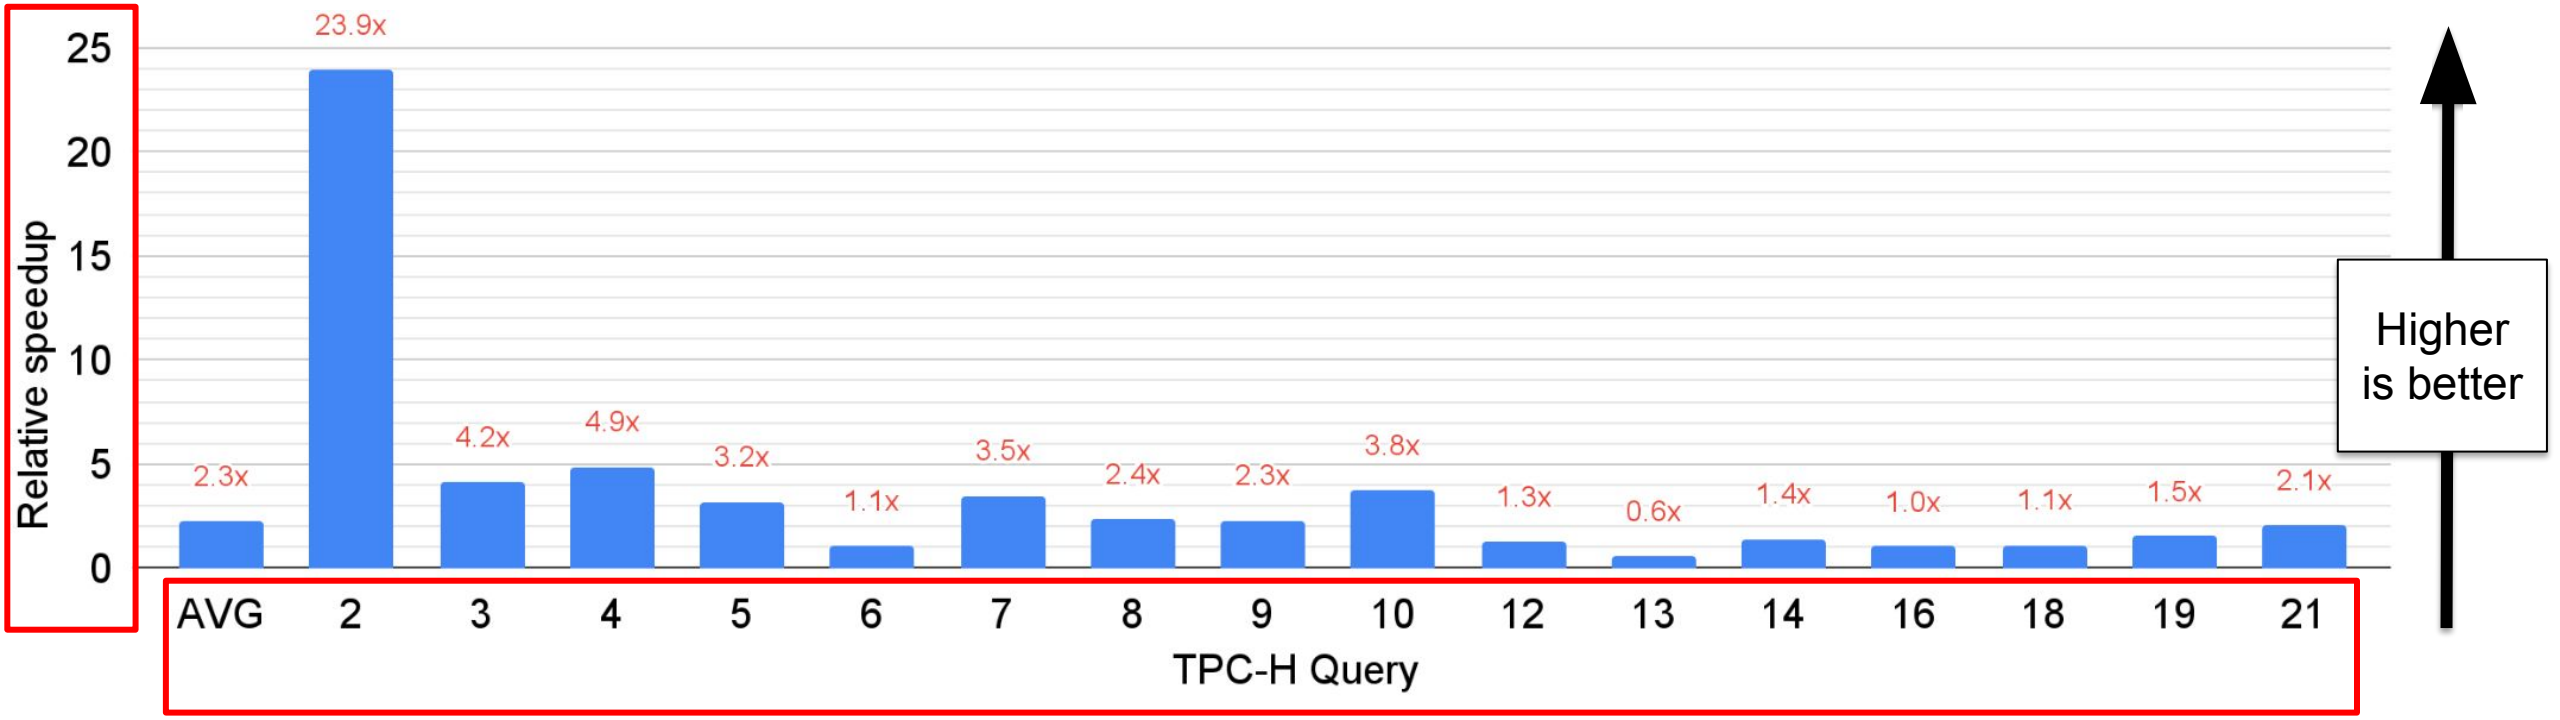

IronSafe is 2.3x faster on average and guarantees strong security and policy compliance

Q2: Policy compliance use cases

GDPR Anti-patterns[1]	Description
Timely deletion	Ensure data is not stored and used indefinitely
Indiscrimination	Control data sharing
Transparency	Transparent sharing of data
Risk agnostic	Prevent untrusted parties from accessing data
Data breaches	Know if data was used by unwanted parties

[1] Supreeth Shastri, Melissa Wasserman, and Vijay Chidambaram. 2021. GDPR anti-patterns. *Commun. ACM* 64, 2 (February 2021)

Summary

- Cloud systems enable high-performance query processing
- Processing queries in a secure and policy compliant manner is challenging
  - Heterogenous TEEs + Policy specification and compliance
- Ironsafe[1]: A secure and policy compliant query processing architecture
  - A heterogenous confidential computing framework
  - IronSafe's declarative policy language
  - A compliance infrastructure for enforcing policies

https://github.com/harshanavkis/ironsafe

Backup slides

Threat model

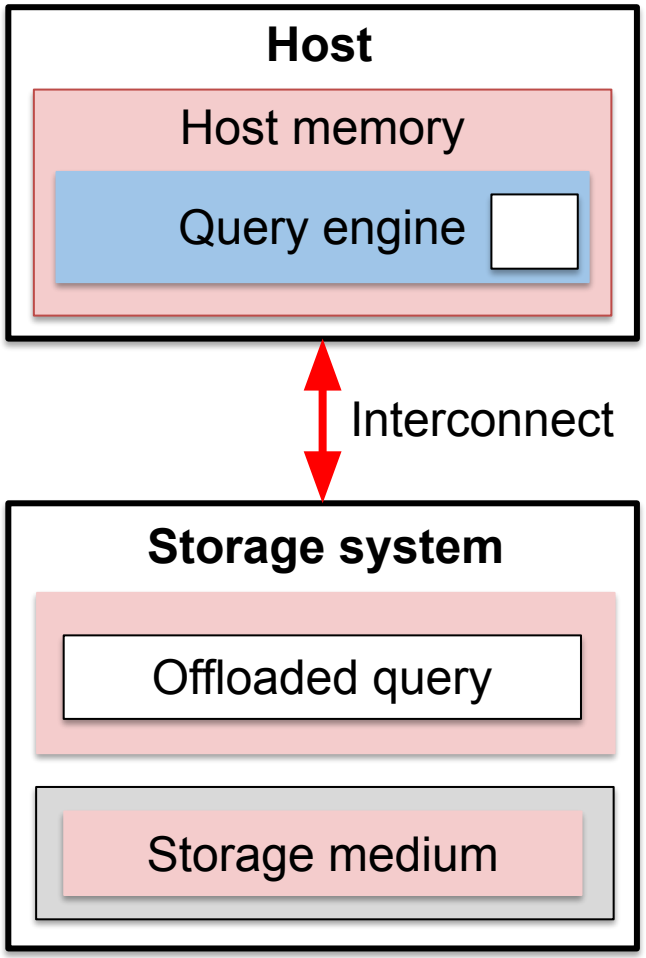

IronSafe overheads

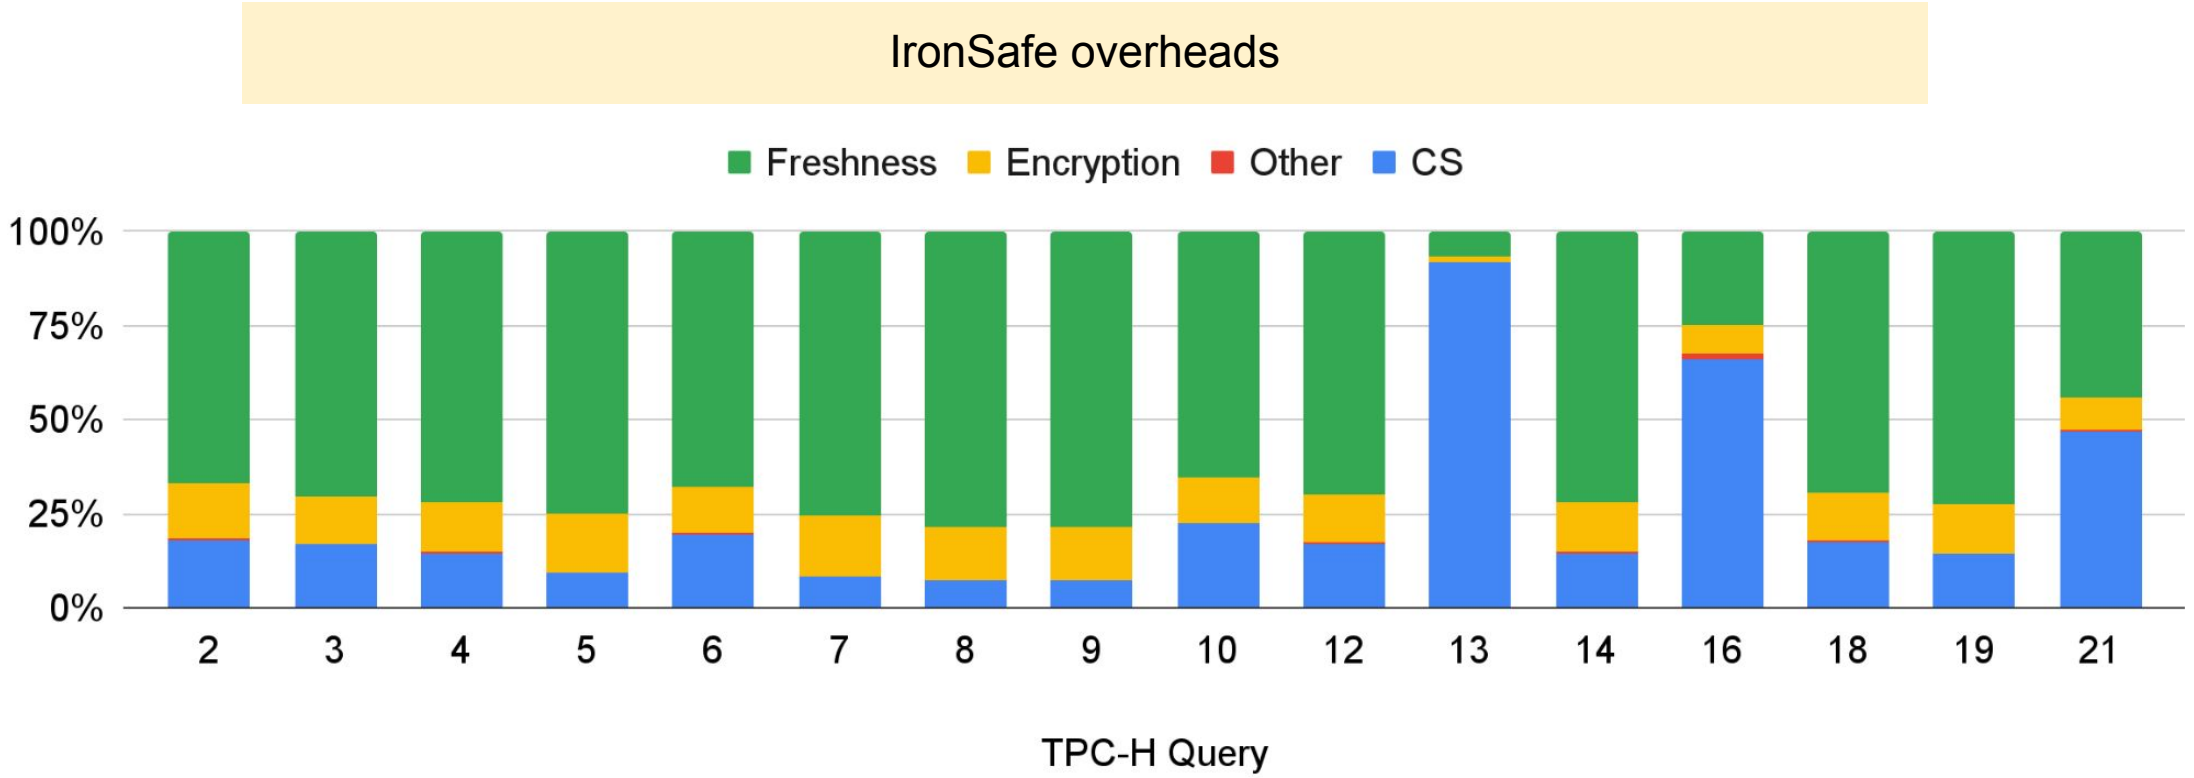

Overheads due to providing confidentiality, integrity and freshness guarantees

Timely deletion of data

read :- **sessionKeyIs**(K_A) | **sessionKeyIs**(K_B) & **le**(T, TIMESTAMP)