UNIVERSITY OF OREGON

Towards Secure Provenance-Based Access Control in Cloud Environments

Adam Bates Ben Mood Masoud Valafar

Kevin Butler

OSRIS Oregon Systems Infrastructure Research & Information Security Laboratory

CODASPY'13, San Antonio, TX, USA
19 February 2013

Data Provenance

UNIVERSITY OF OREGON

Def: prov·en·ance \prä-və-nän(t)s\ *n*:

- A metadata history detailing an object's derivation.
- Provides context needed to answer questions like:
  - *What applications operated on this data?*
  - *What datasets helped produce this data?*
  - *In what environment was this data processed?*

Cloud Provenance Uses

UNIVERSITY OF OREGON

- Regulatory compliance.
- Debug experimental results.*
- Detect and avoid faulty data propagation.*
- Improve text search results.*
- Digital attack forensics.†

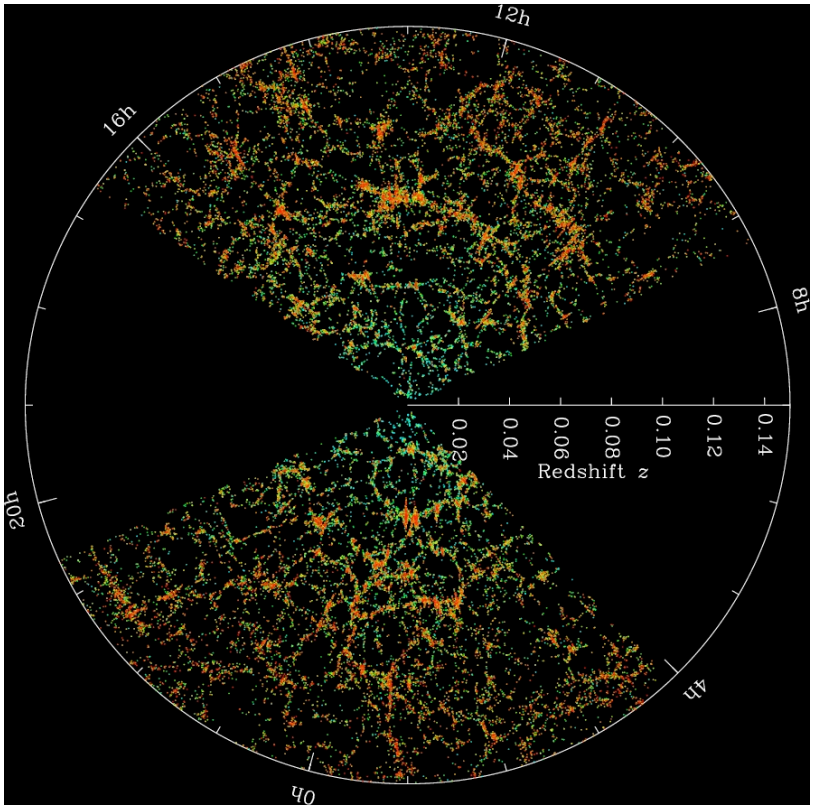

*Data processing in the cloud benefits from the added context of data provenance.**

* Muniswamy-Reddy et al, FAST'10

† Galante et al. "Sony Network Breach..."

Obstacles

UNIVERSITY OF OREGON

Challenges between us and our provenance-aware cloud:

- Host-level collection
- Storage
- Distributed security
- Distributed management
- Killer apps

Obstacles

UNIVERSITY OF OREGON

Challenges between us and our provenance-aware cloud:

- Host-level collection
- Storage
- Distributed security
- Distributed management
- Killer apps

Overview

UNIVERSITY OF OREGON

This work introduces:

- System and protocols to secure and manage provenance sent to cloud.
- Proof-of-concept provenance-aware cloud access control mechanisms.
- Performance evaluation that demonstrates minimal imposed overhead (~14%).

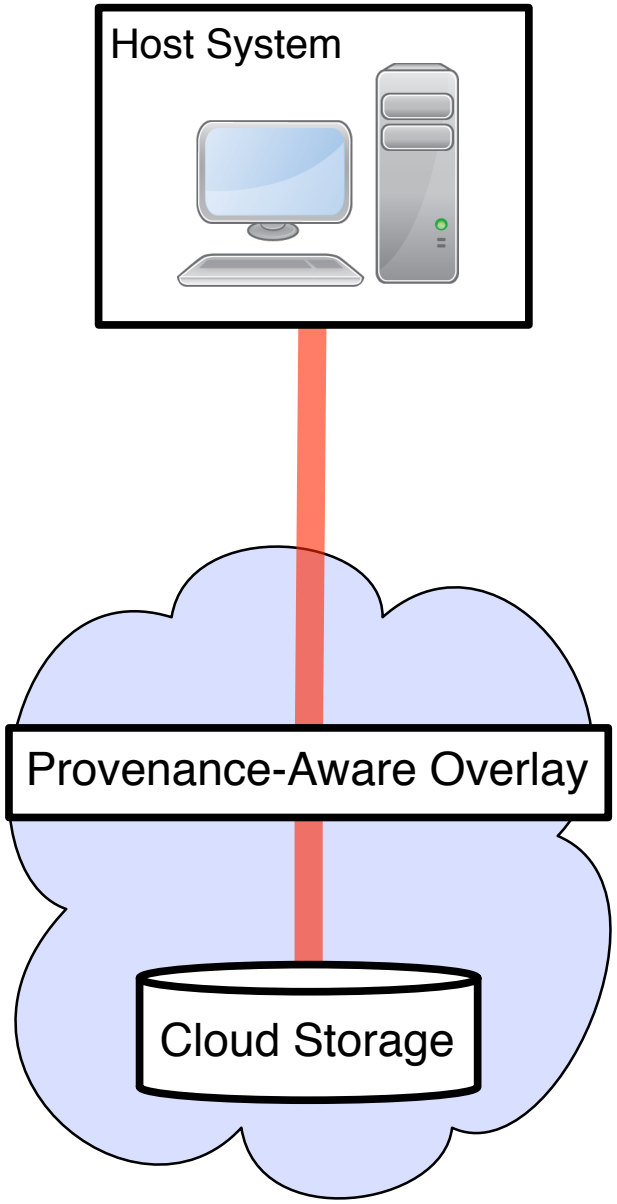

Cloud Provenance Authority

UNIVERSITY OF OREGON

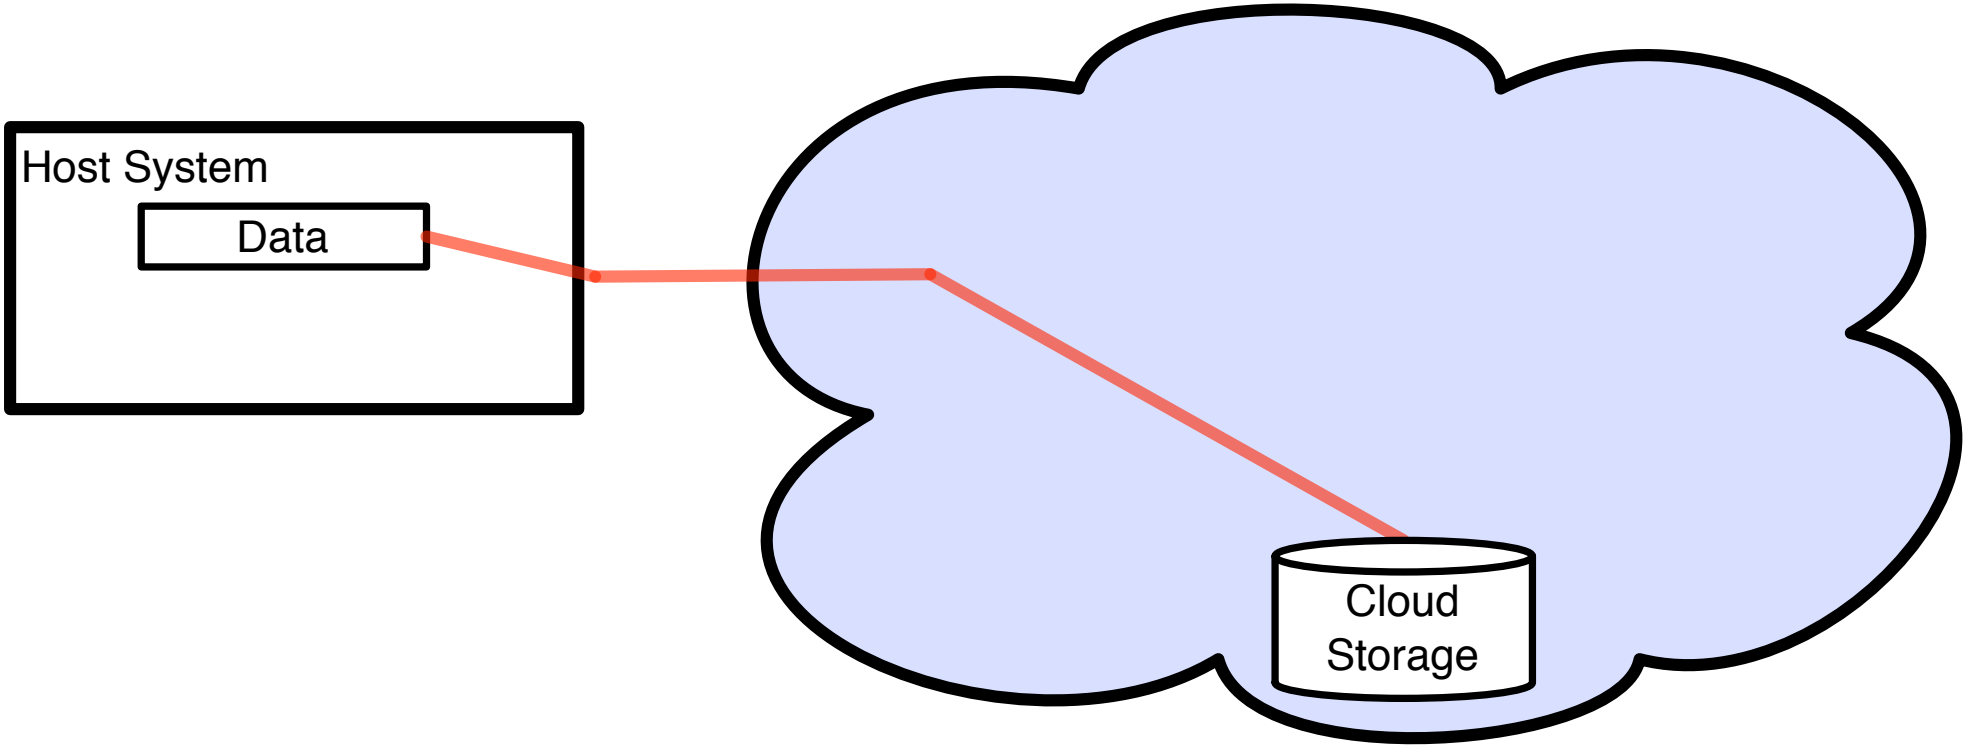

Cloud Provenance Authority

UNIVERSITY OF OREGON

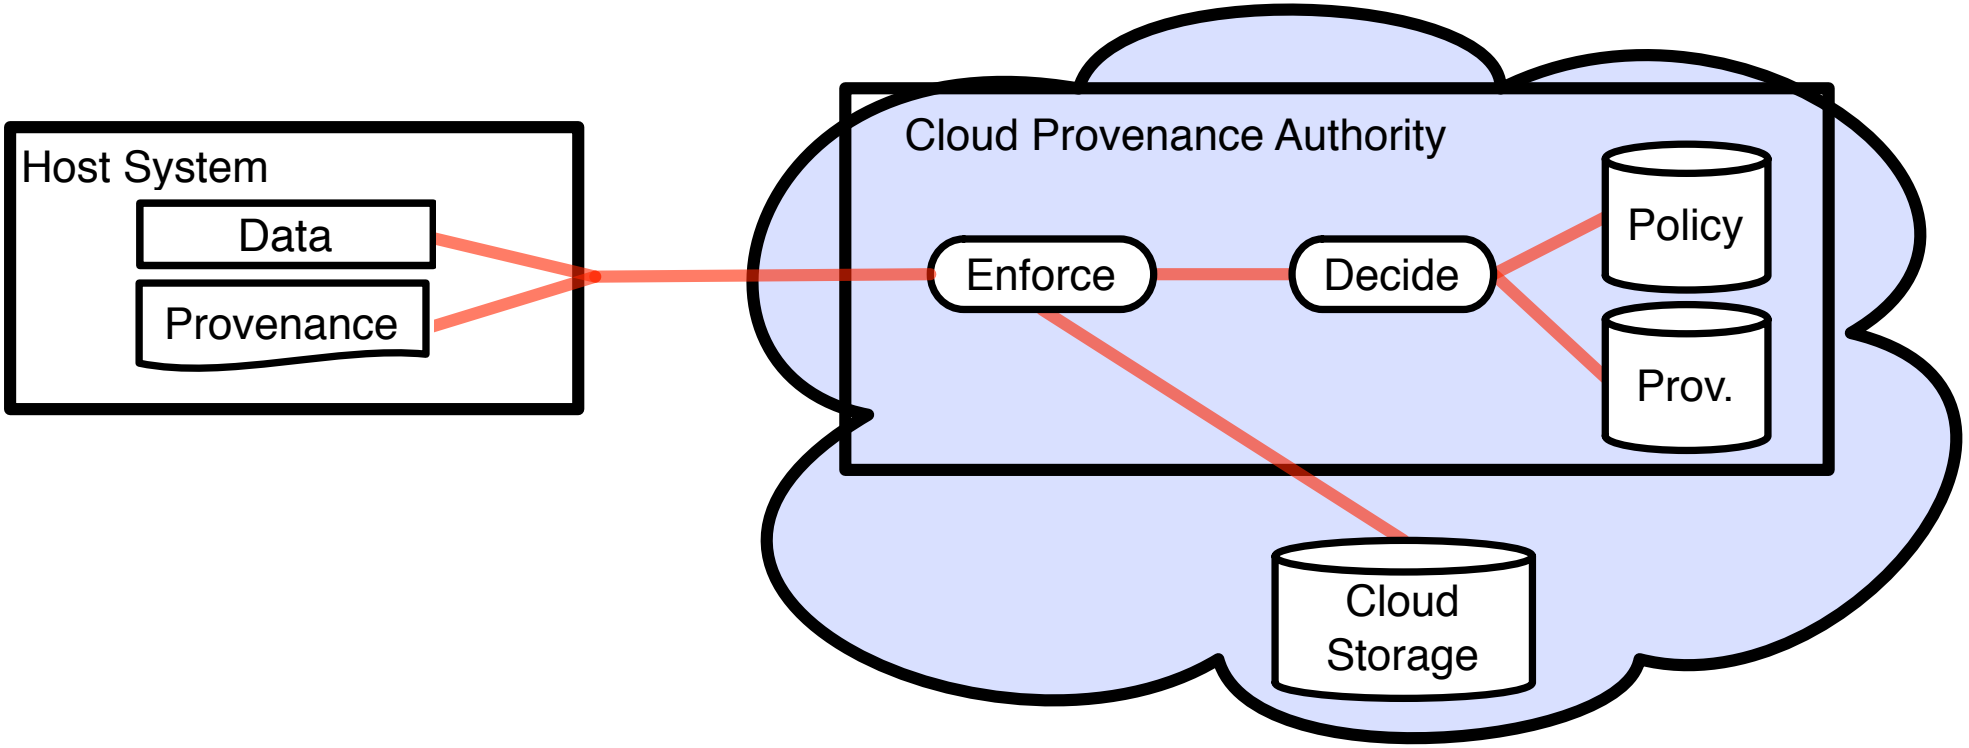

Enforcement Points: Interact with clients and mediate access to cloud storage.

Decision Points: Issues access decisions and stores provenance and security policies.

Cloud Provenance Authority

UNIVERSITY OF OREGON

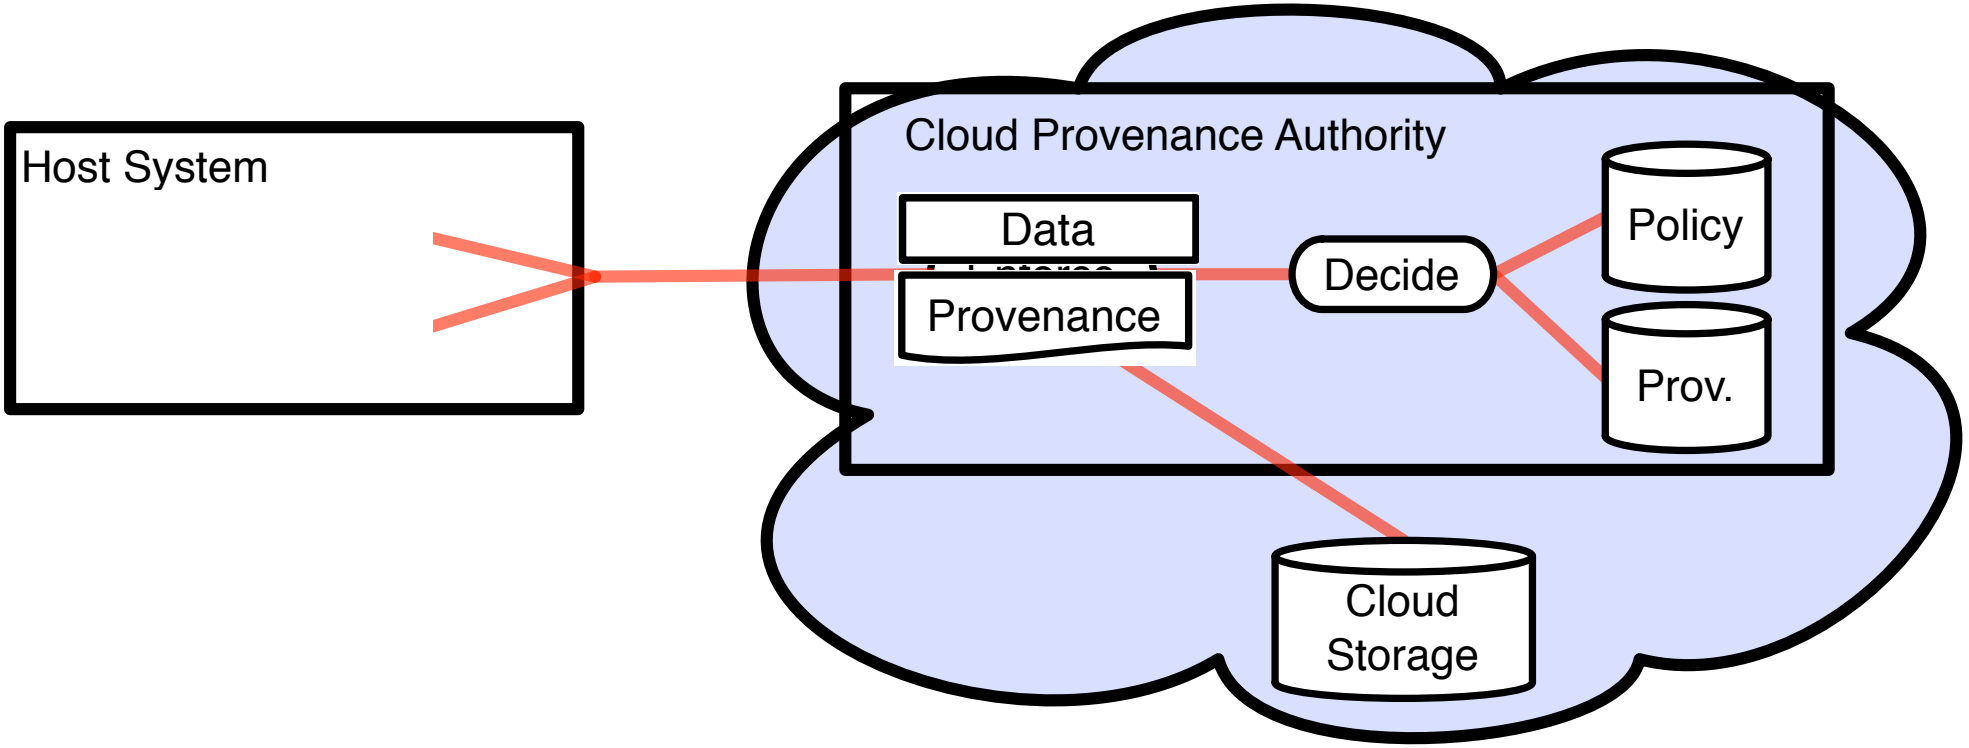

Enforcement Points: Interact with clients and mediate access to cloud storage.

Decision Points: Issues access decisions and stores provenance and security policies.

Cloud Provenance Authority

UNIVERSITY OF OREGON

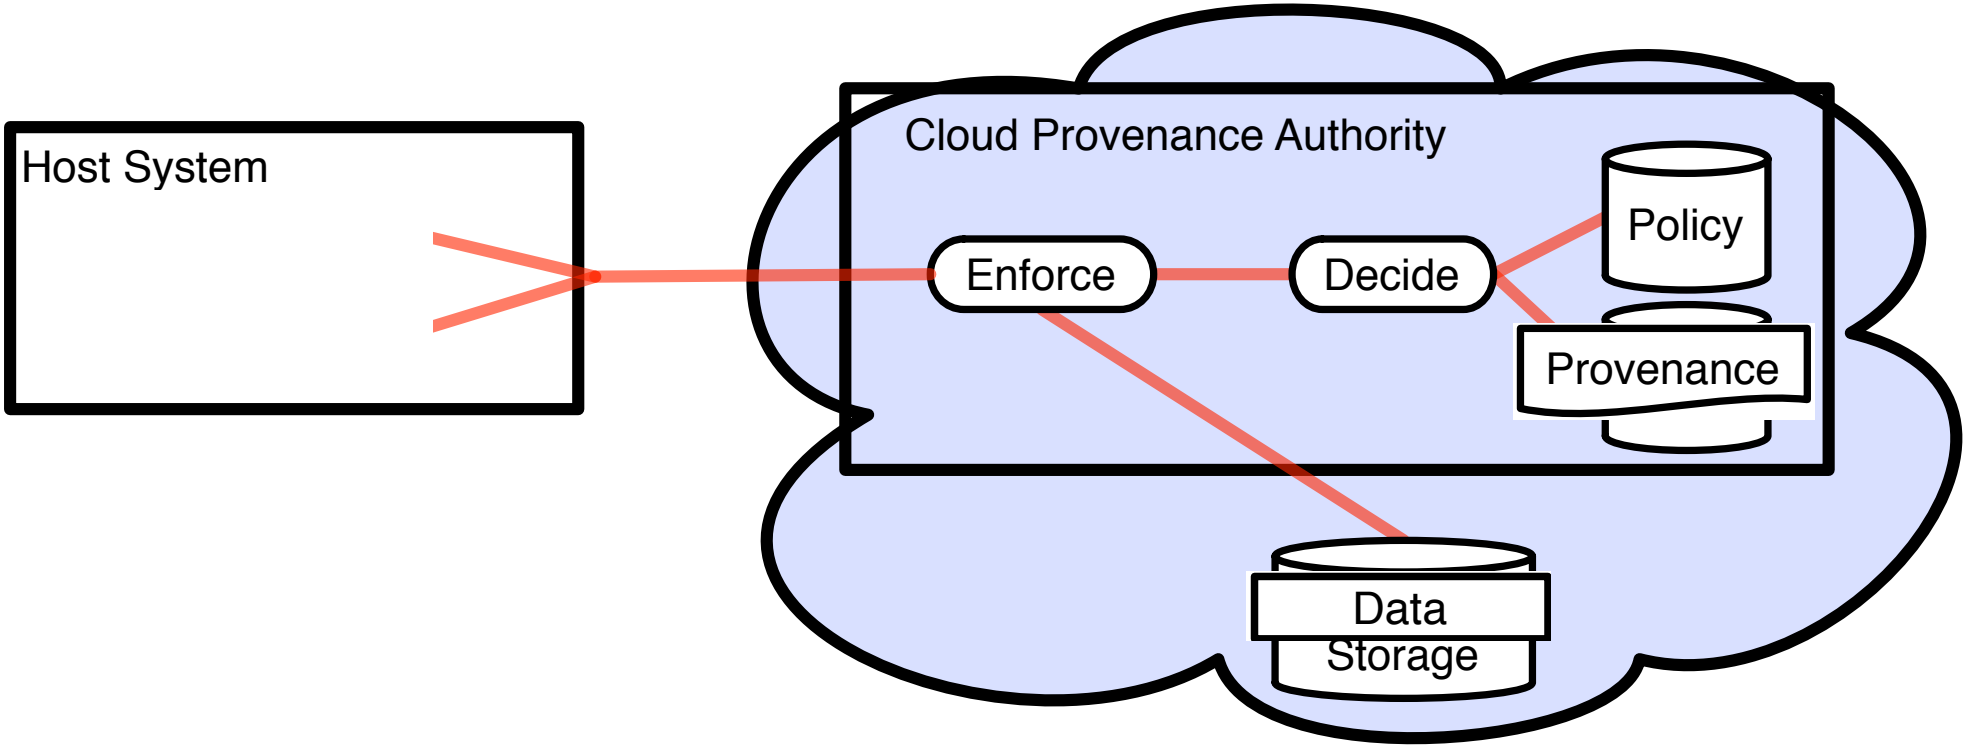

Enforcement Points: Interact with clients and mediate access to cloud storage.

Decision Points: Issues access decisions and stores provenance and security policies.

Commitment Protocol

UNIVERSITY OF OREGON

Client (C)

Enforcement Point

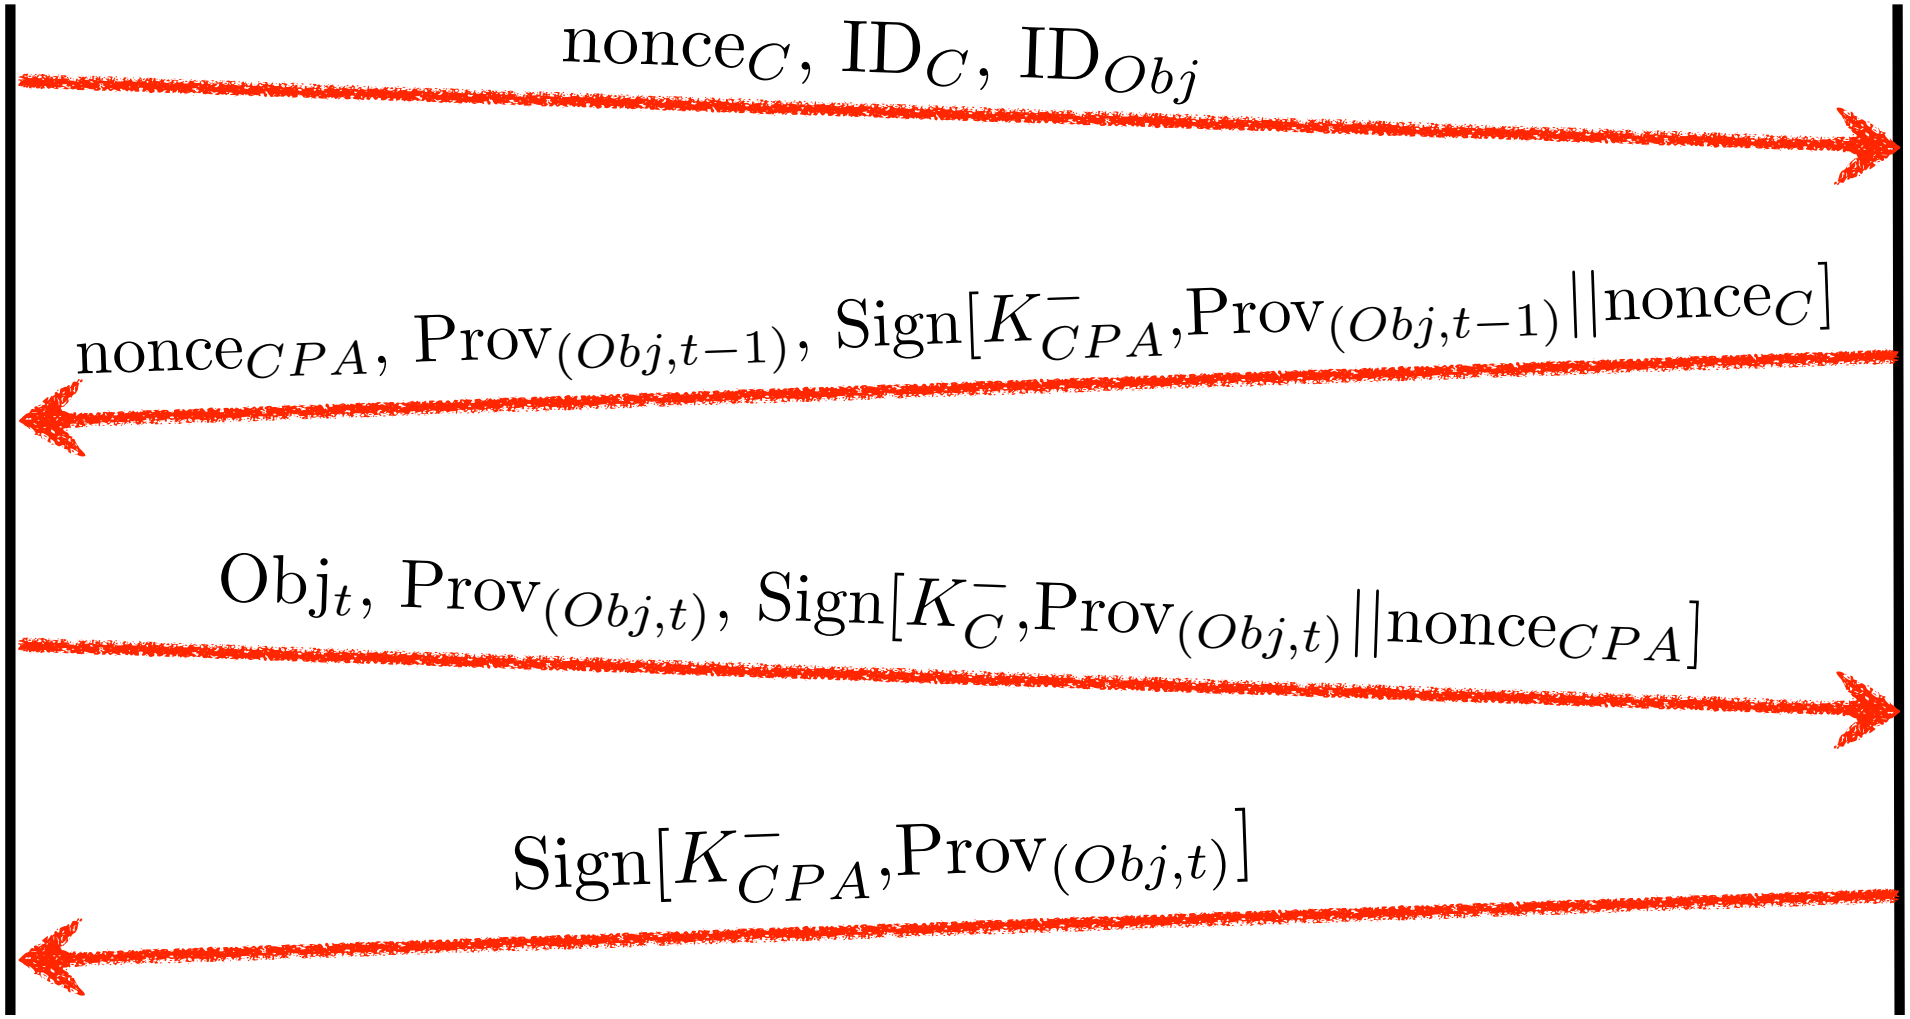

Commitment Protocol

UNIVERSITY OF OREGON

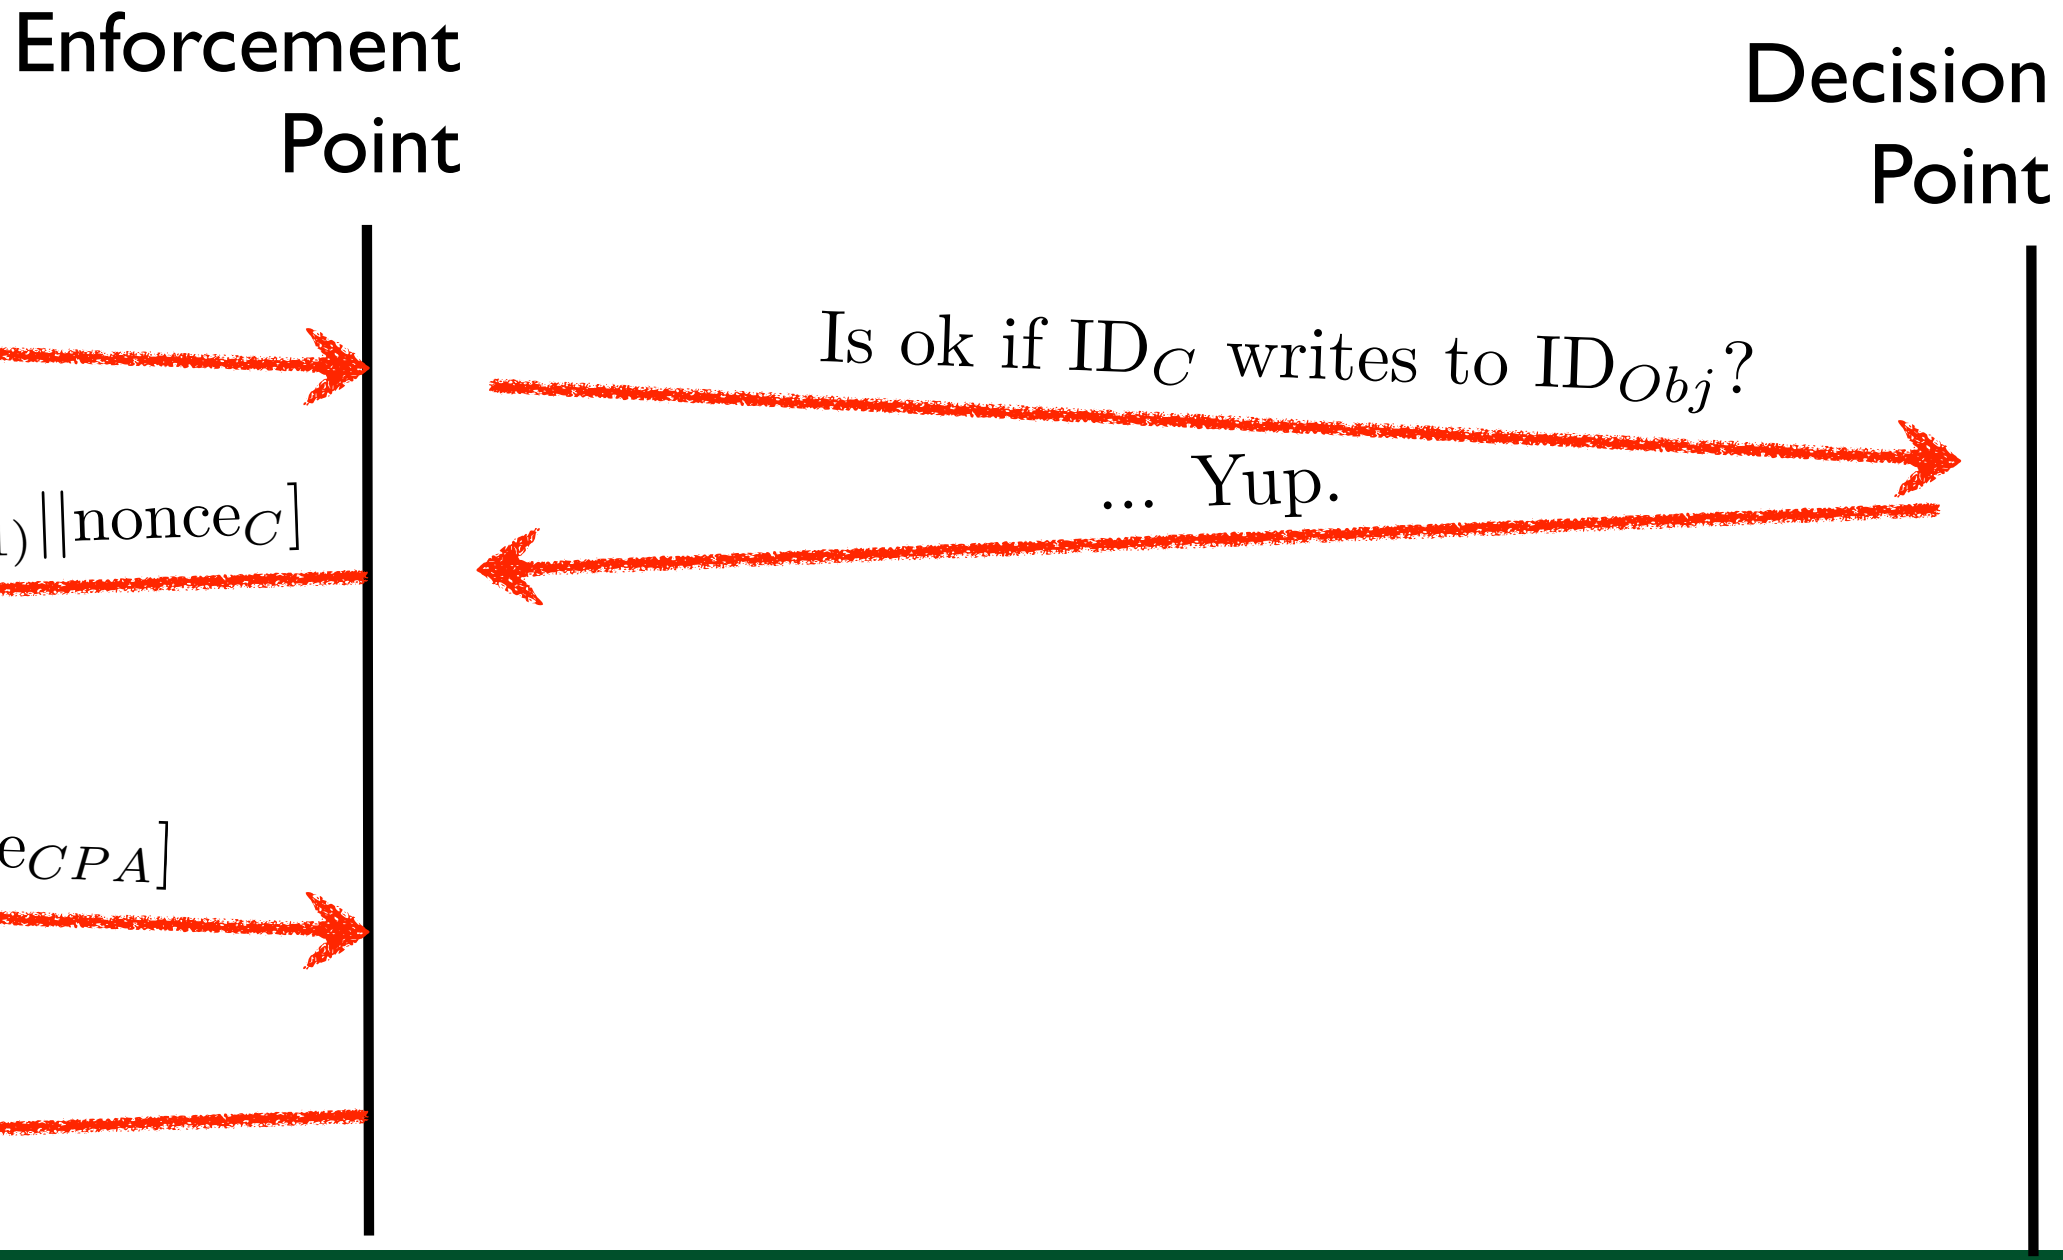

Commitment Protocol

UNIVERSITY OF OREGON

Client (C)

Enforcement Point

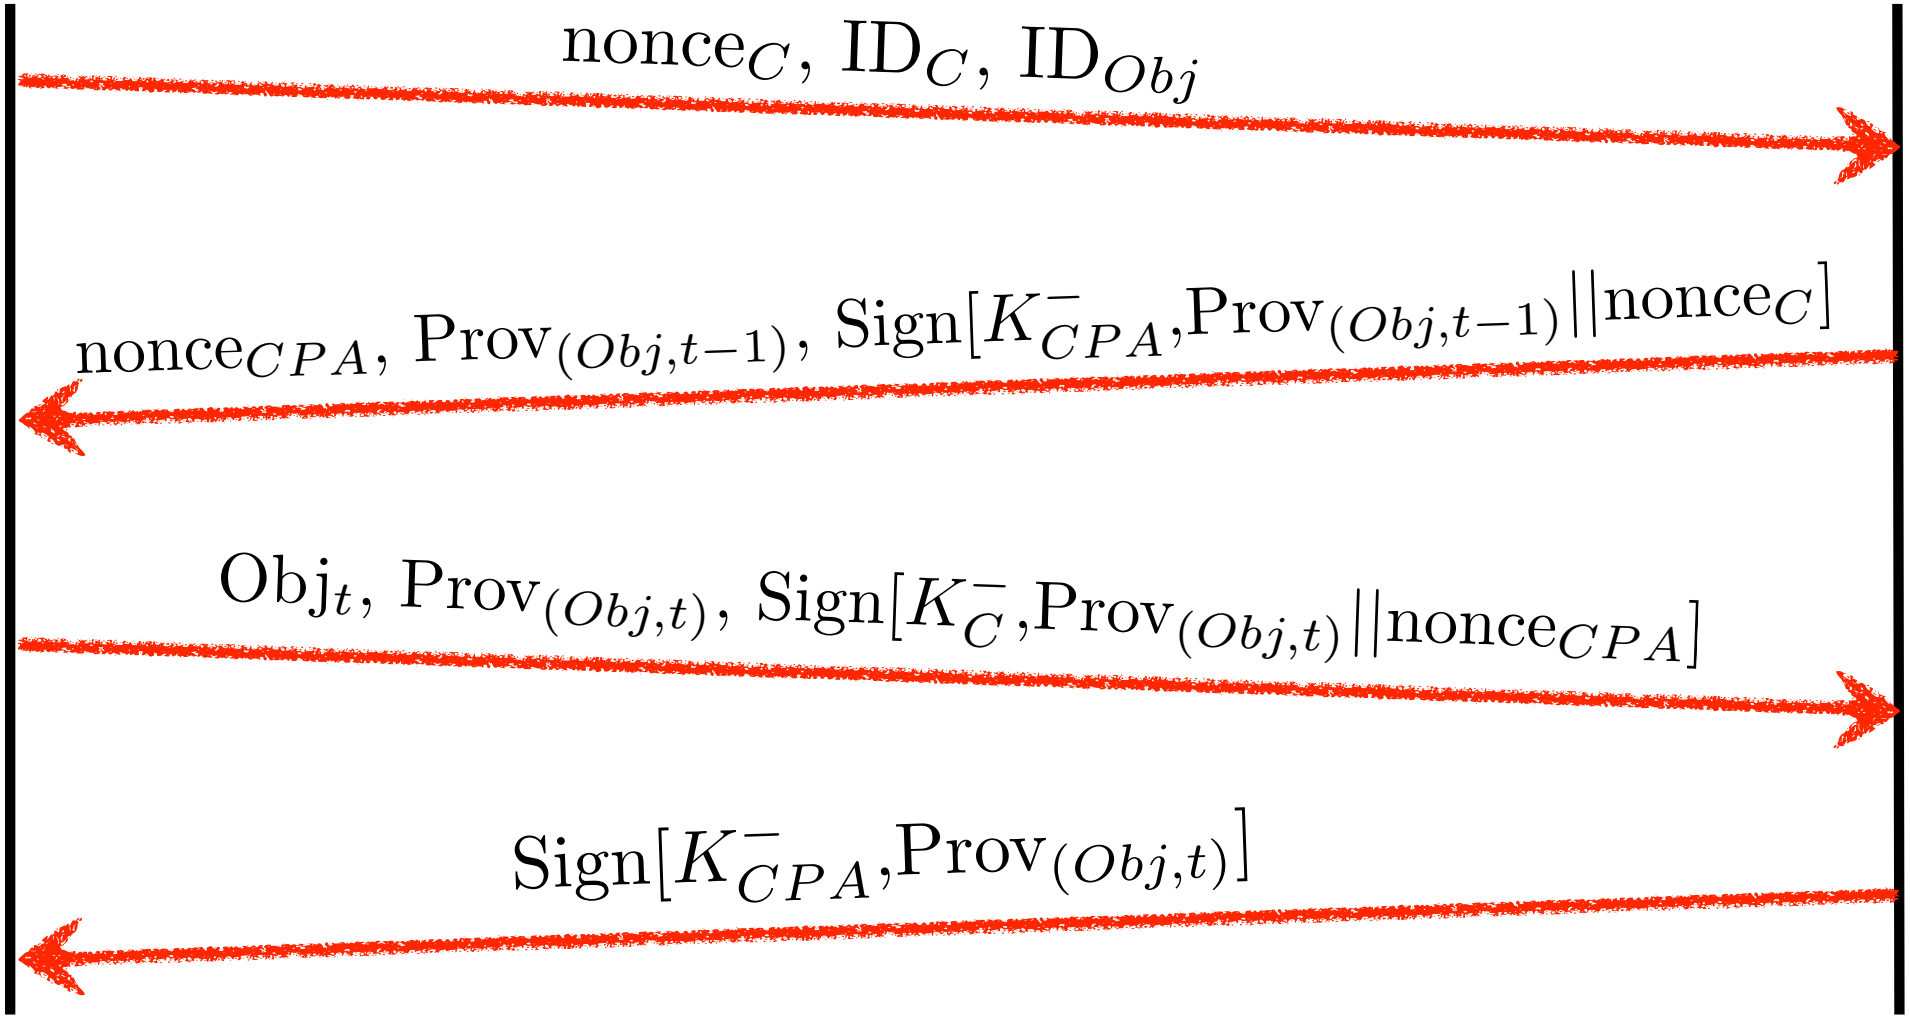

Provenance Chains

UNIVERSITY OF OREGON

Client (C)

Cloud Provenance Authority (CPA)

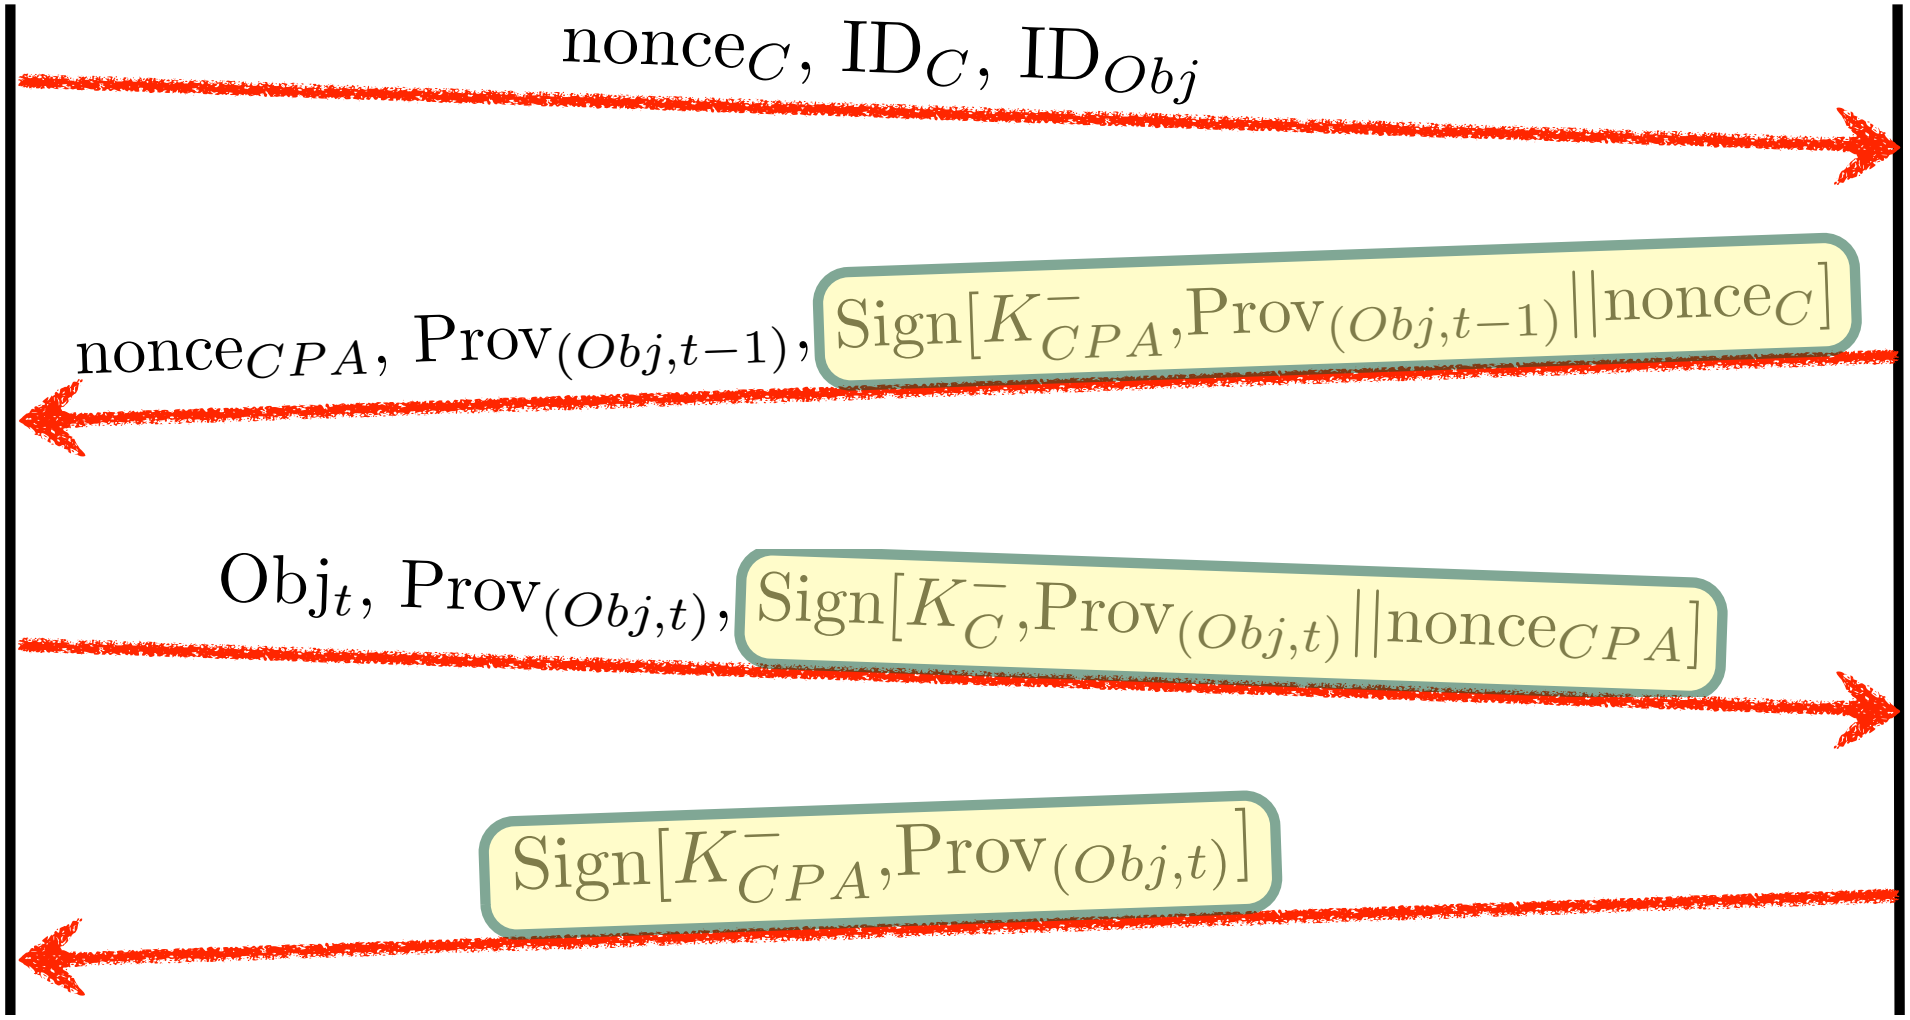

Management & Retrieval

UNIVERSITY OF OREGON

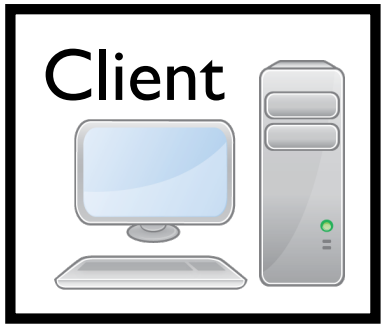

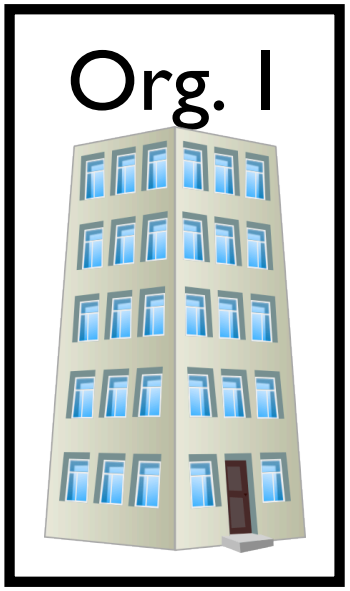

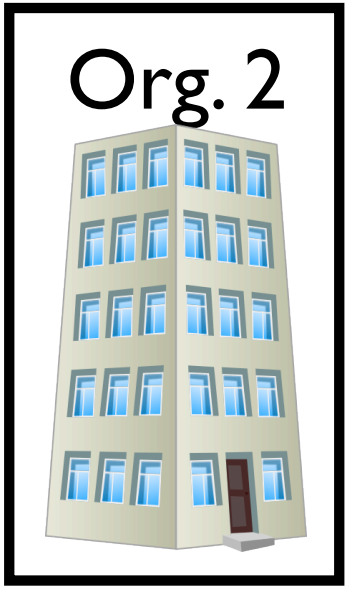

Management & Retrieval

UNIVERSITY OF OREGON

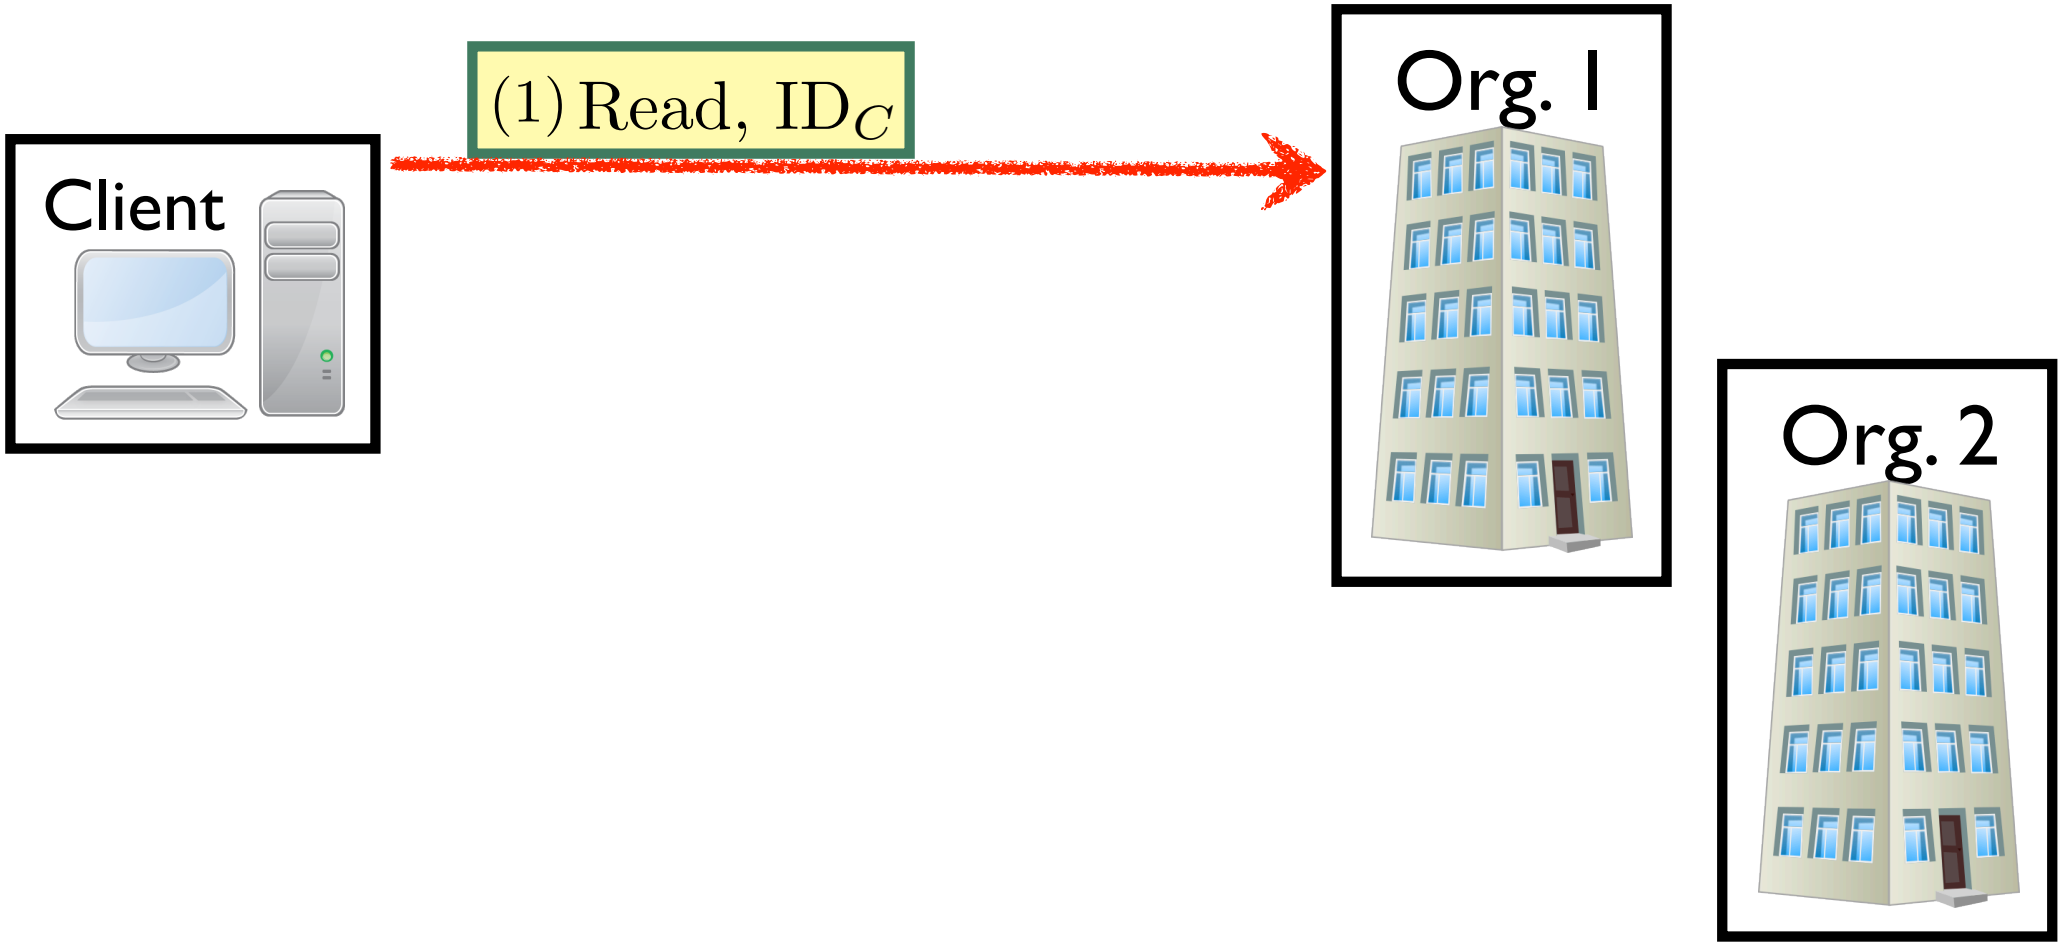

Management & Retrieval

UNIVERSITY OF OREGON

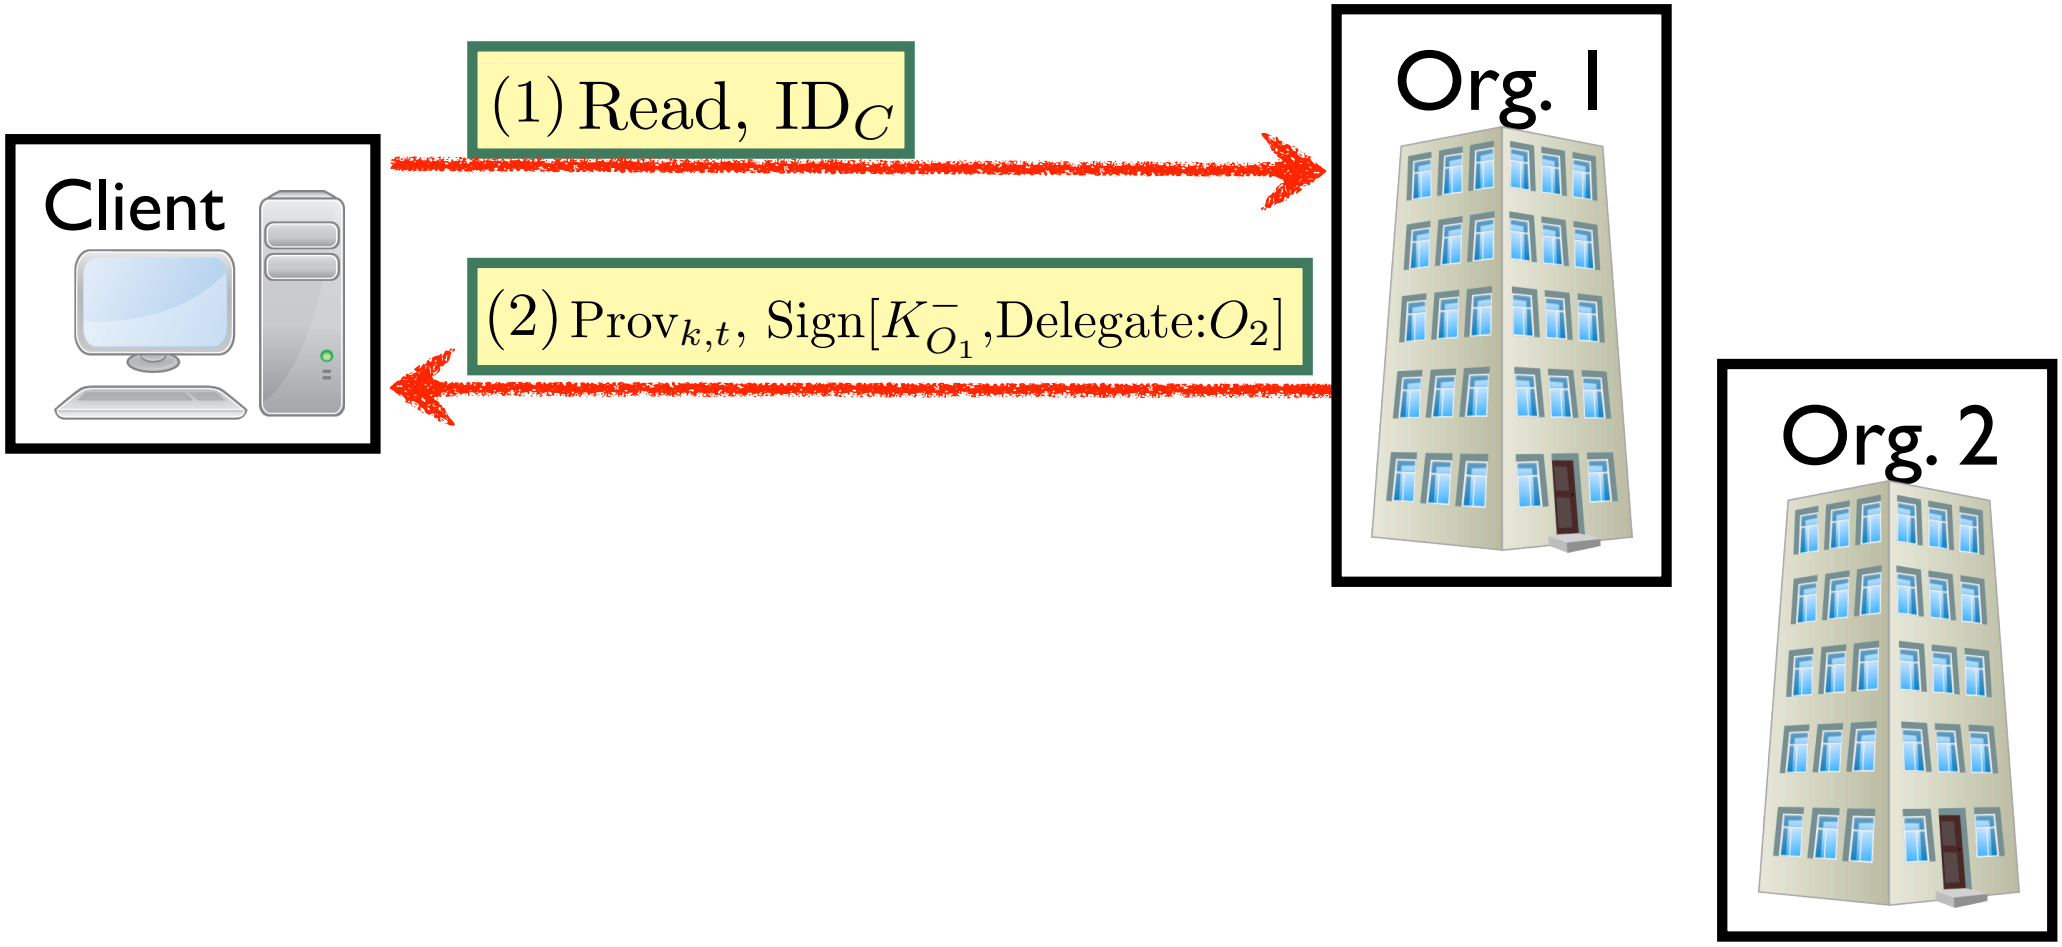

Management & Retrieval

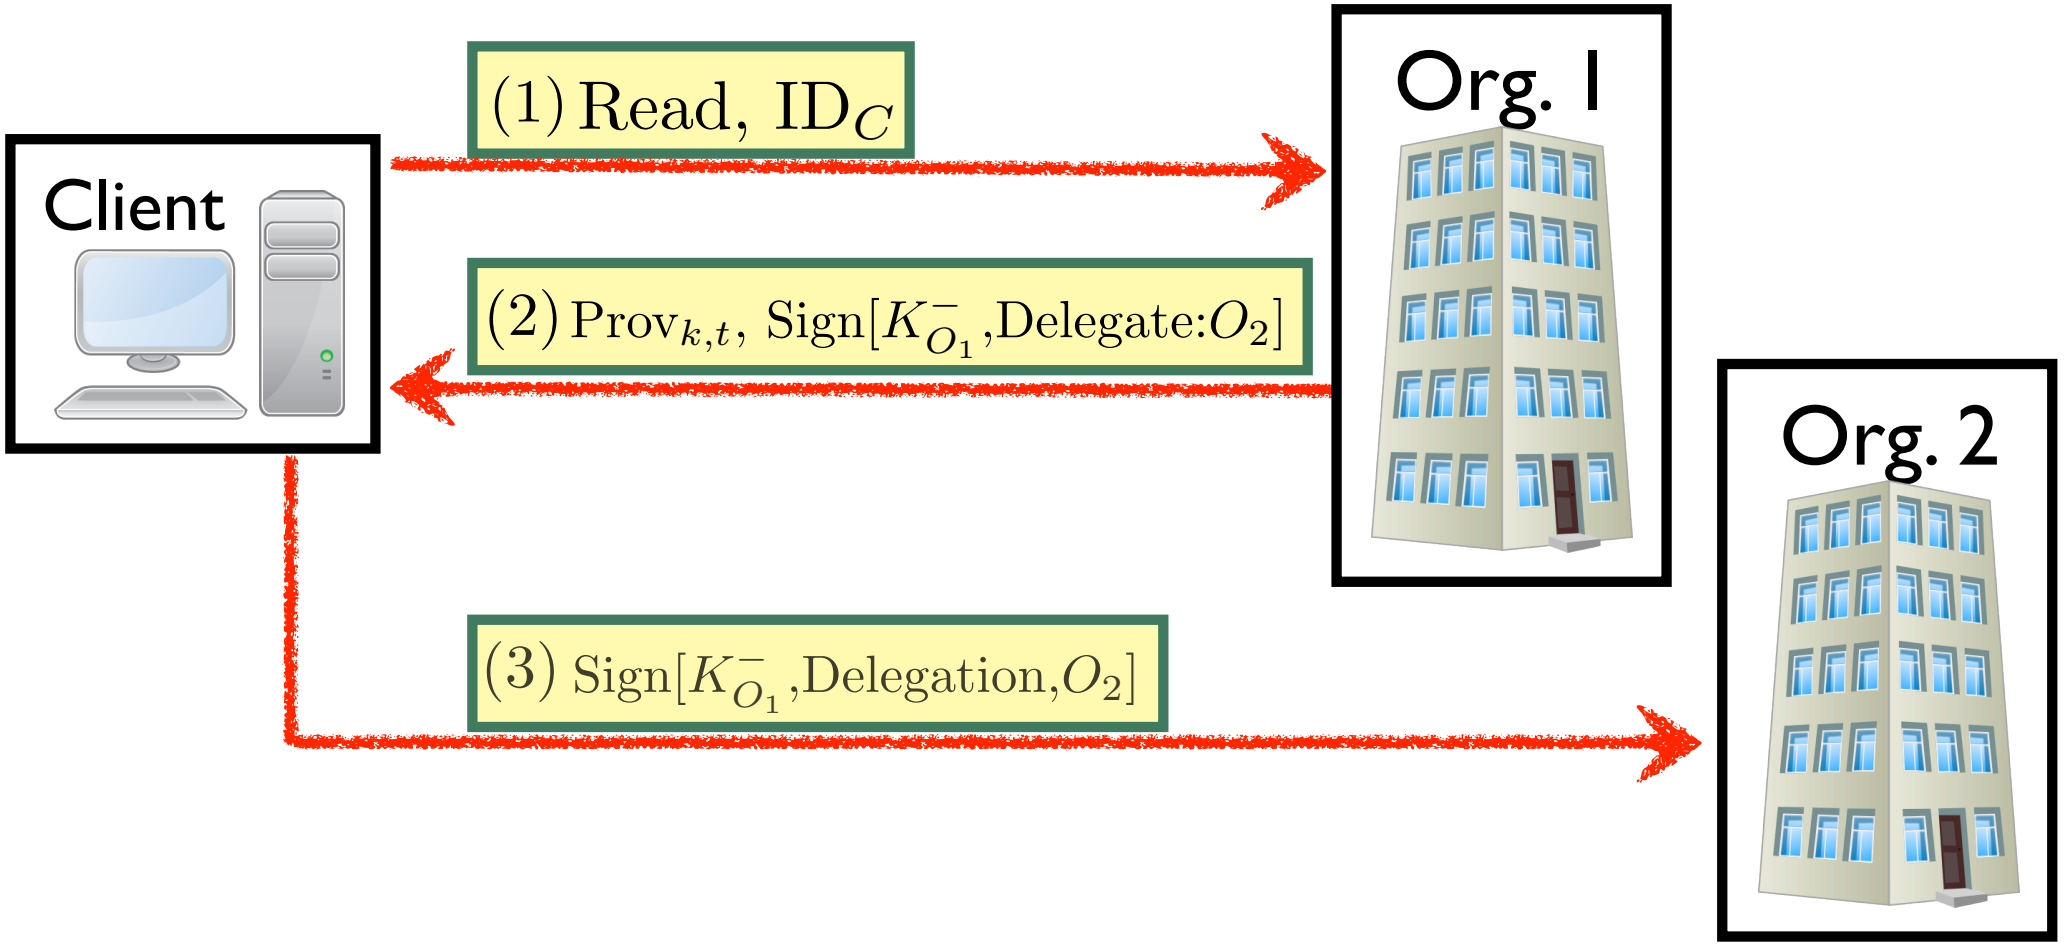

Management & Retrieval

UNIVERSITY OF OREGON

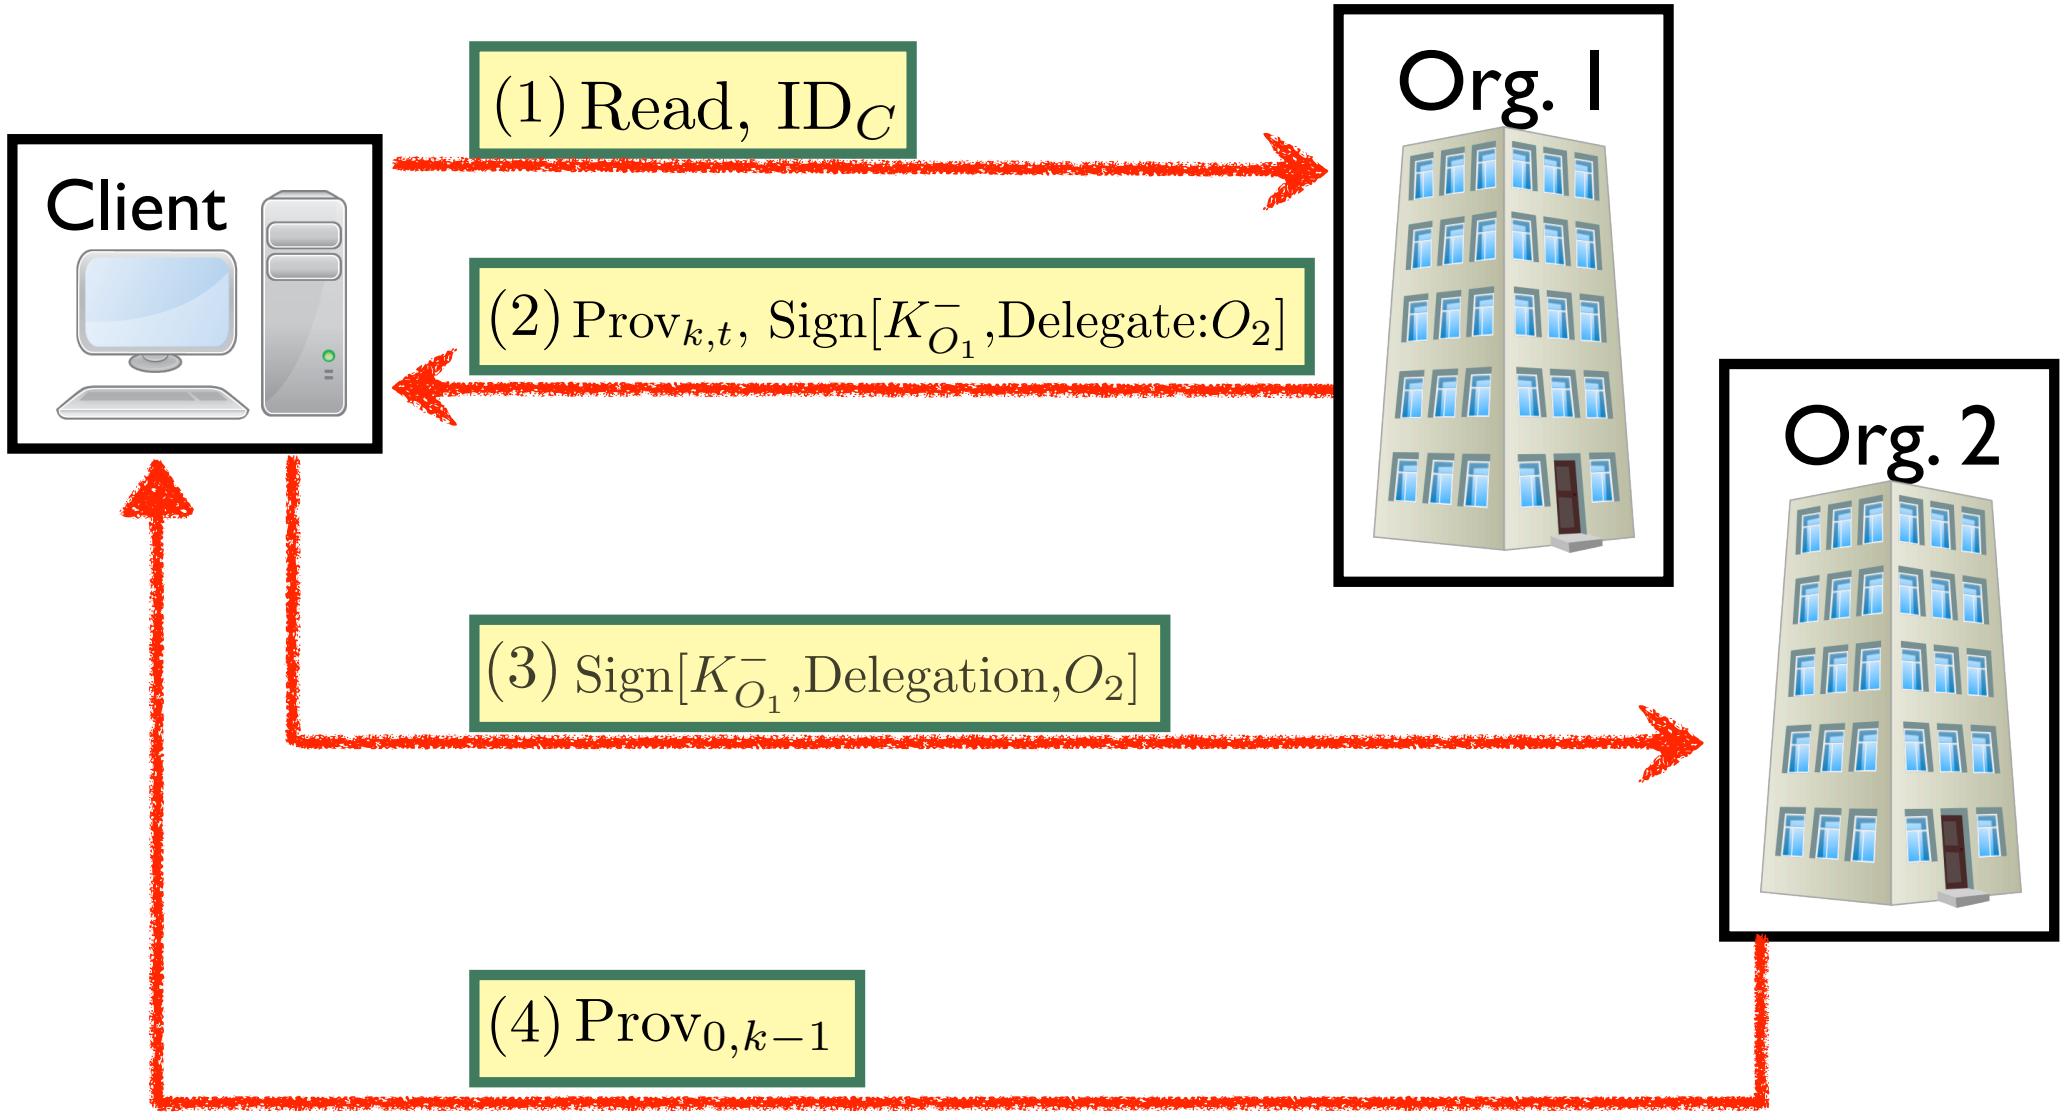

Access Control Example

UNIVERSITY OF OREGON

Client attempts to write 2.1 to Amazon AWS EU (Ireland)...

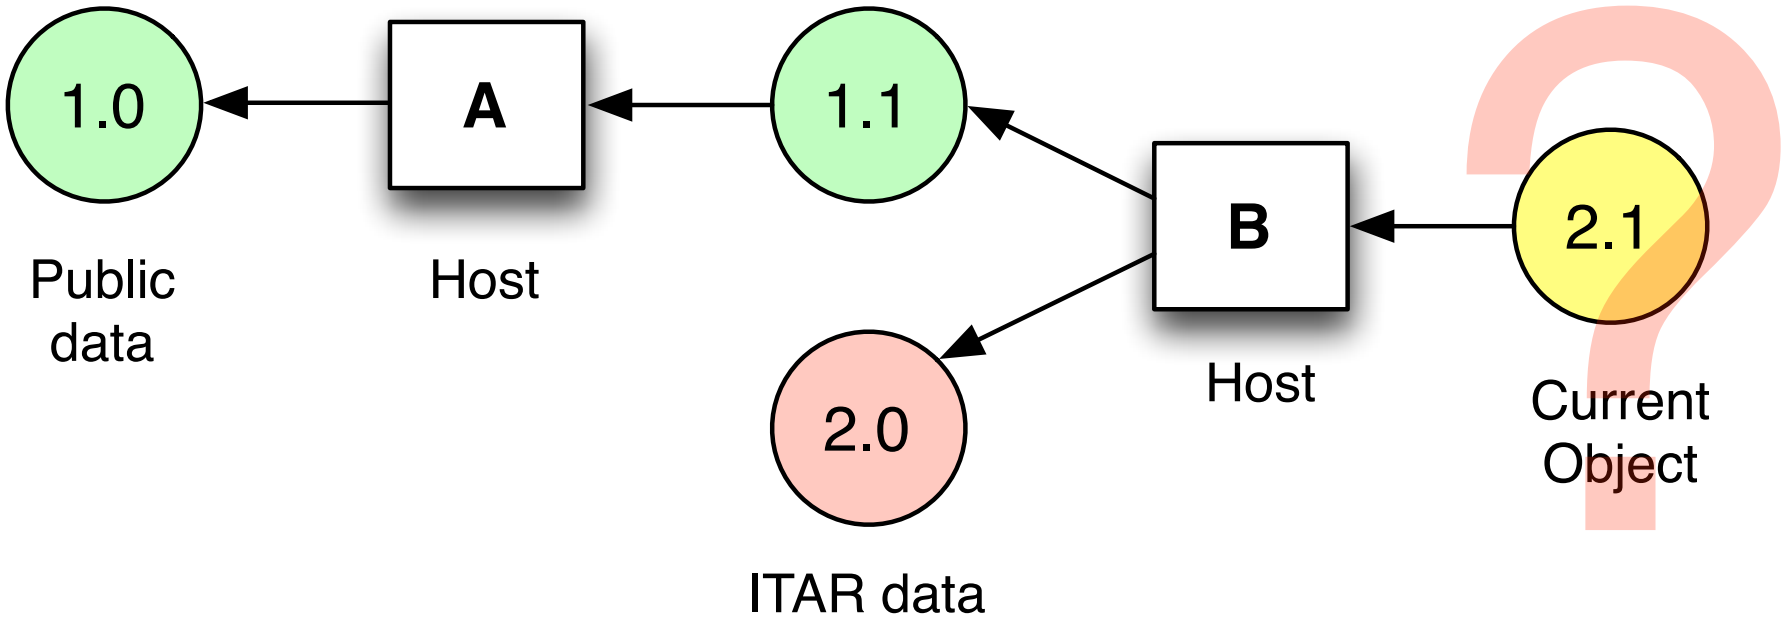

Access Control Example

UNIVERSITY OF OREGON

Client attempts to write 2.1 to Amazon AWS EU (Ireland)...

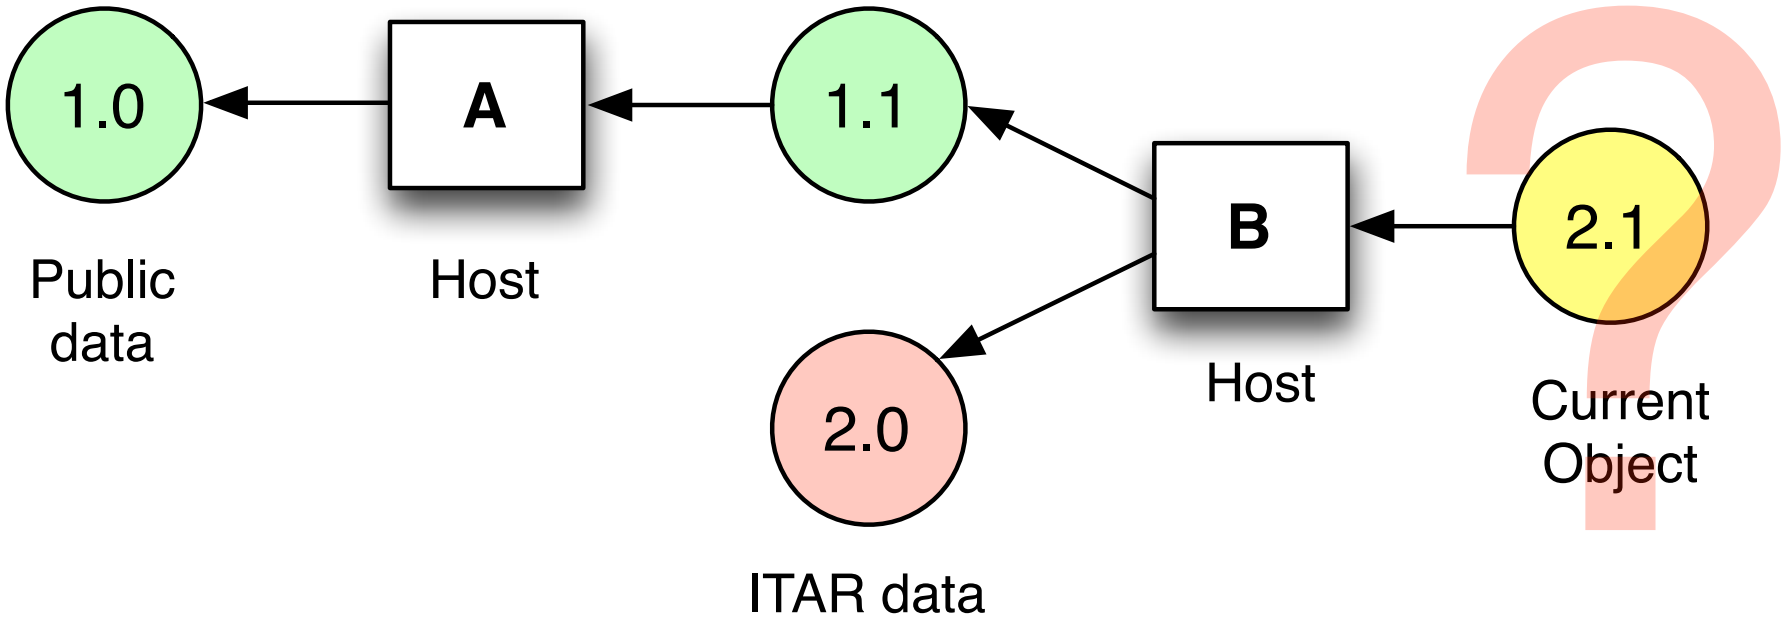

Write Request is denied at the Policy Decision Point because Object 2.1 is derived from ITAR data (Object 2.0)!

Implementation

UNIVERSITY OF OREGON

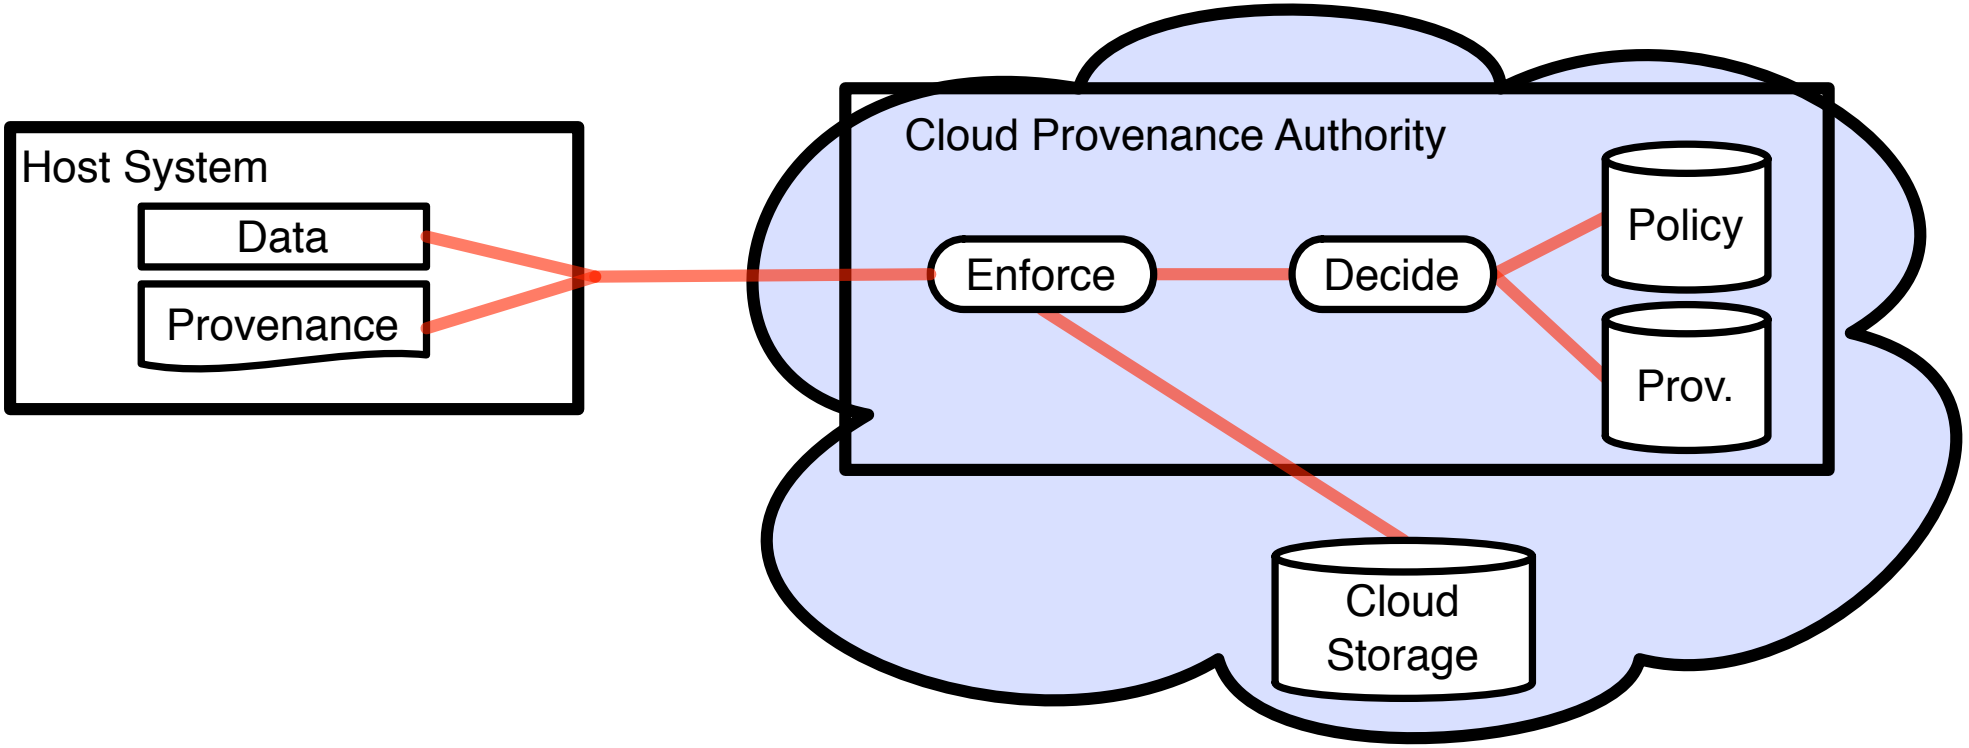

Cloud: University of Oregon ACISS OpenStack KVM.

Components: VMs with 2vCPUs, 4 GB memory.

Communication: Amazon S3 REST API.

Performance

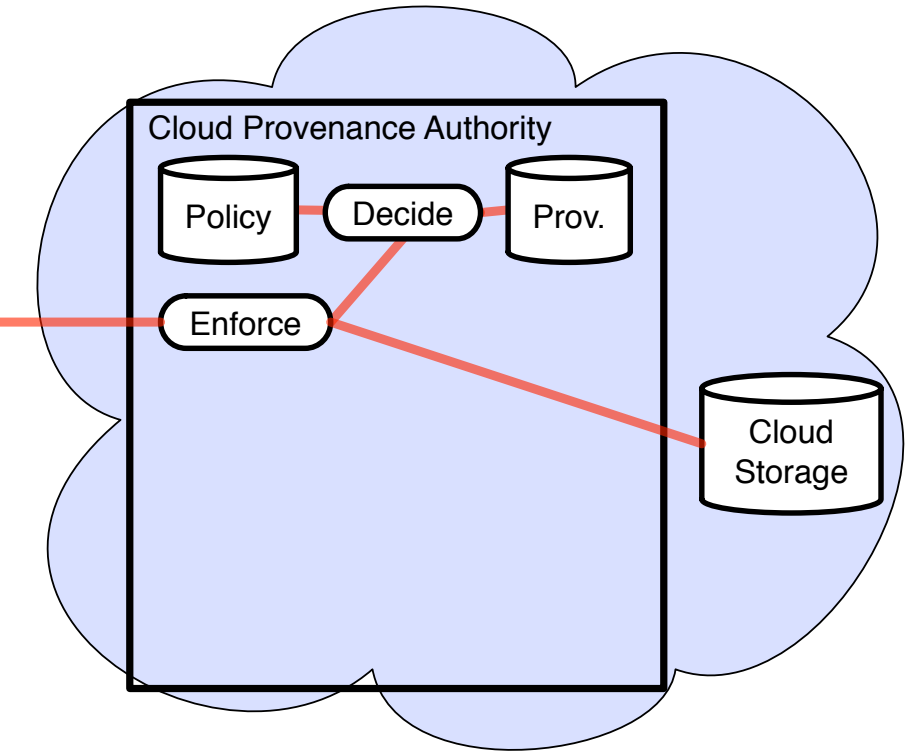

- Major overhead imposed by redundant data transmission (Client to PEP, PEP to Storage).

Performance

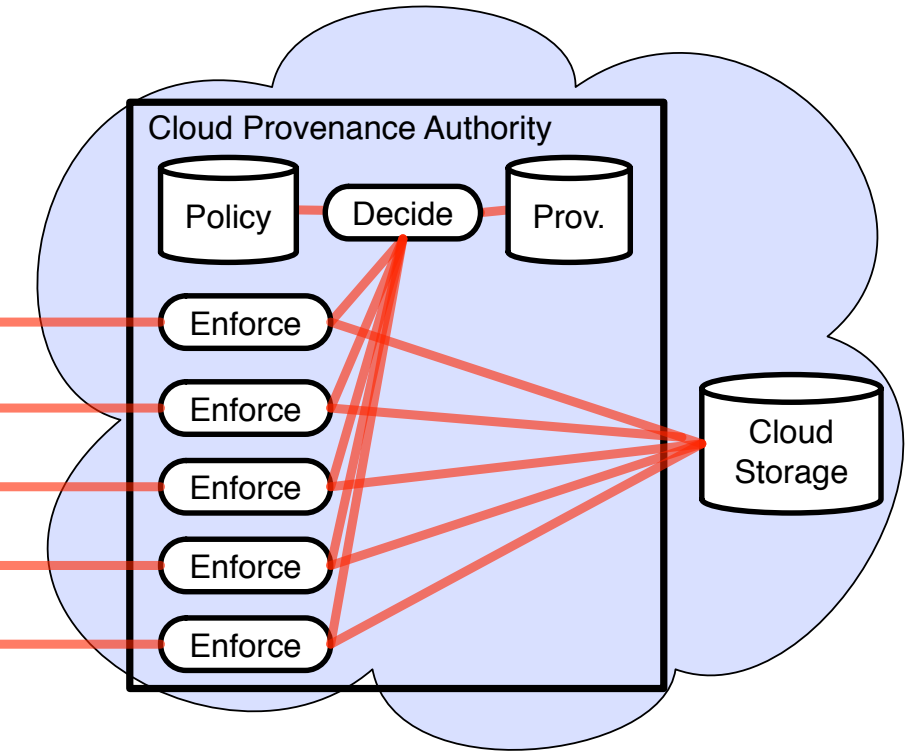

- Major overhead imposed by redundant data transmission (Client to PEP, PEP to Storage).
- *By distributing the PEP workload,* ***we reduced overhead to just 14%.***

Performance

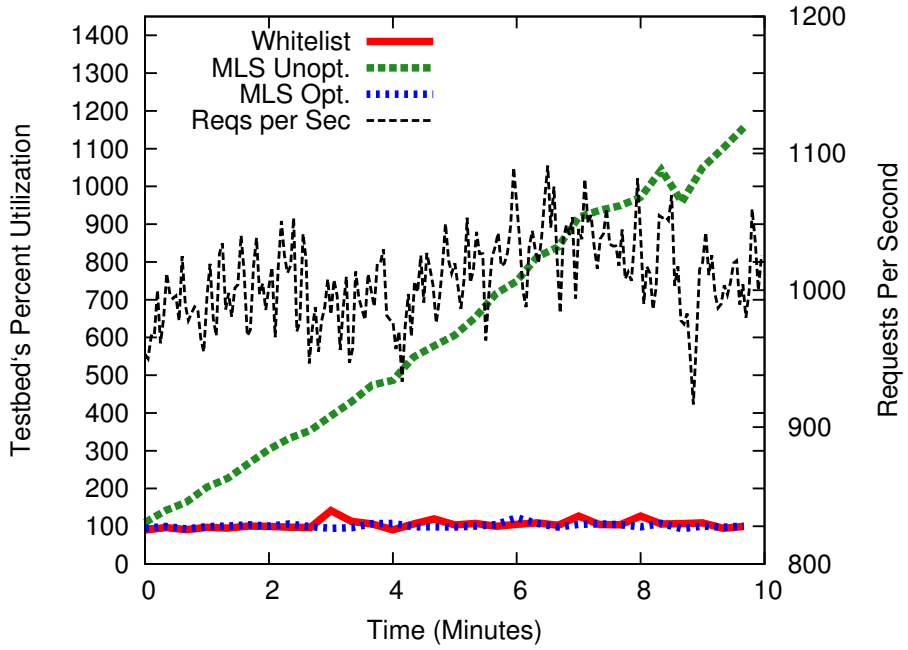

- Major overhead imposed by redundant data transmission (Client to PEP, PEP to Storage).
- *By distributing the PEP workload,* ***we reduced overhead to just 14%.***
- *Under realistic server workloads,* ***our access control mechanism handled 1000 requests per second.***

Performance: Access Control

UNIVERSITY OF OREGON

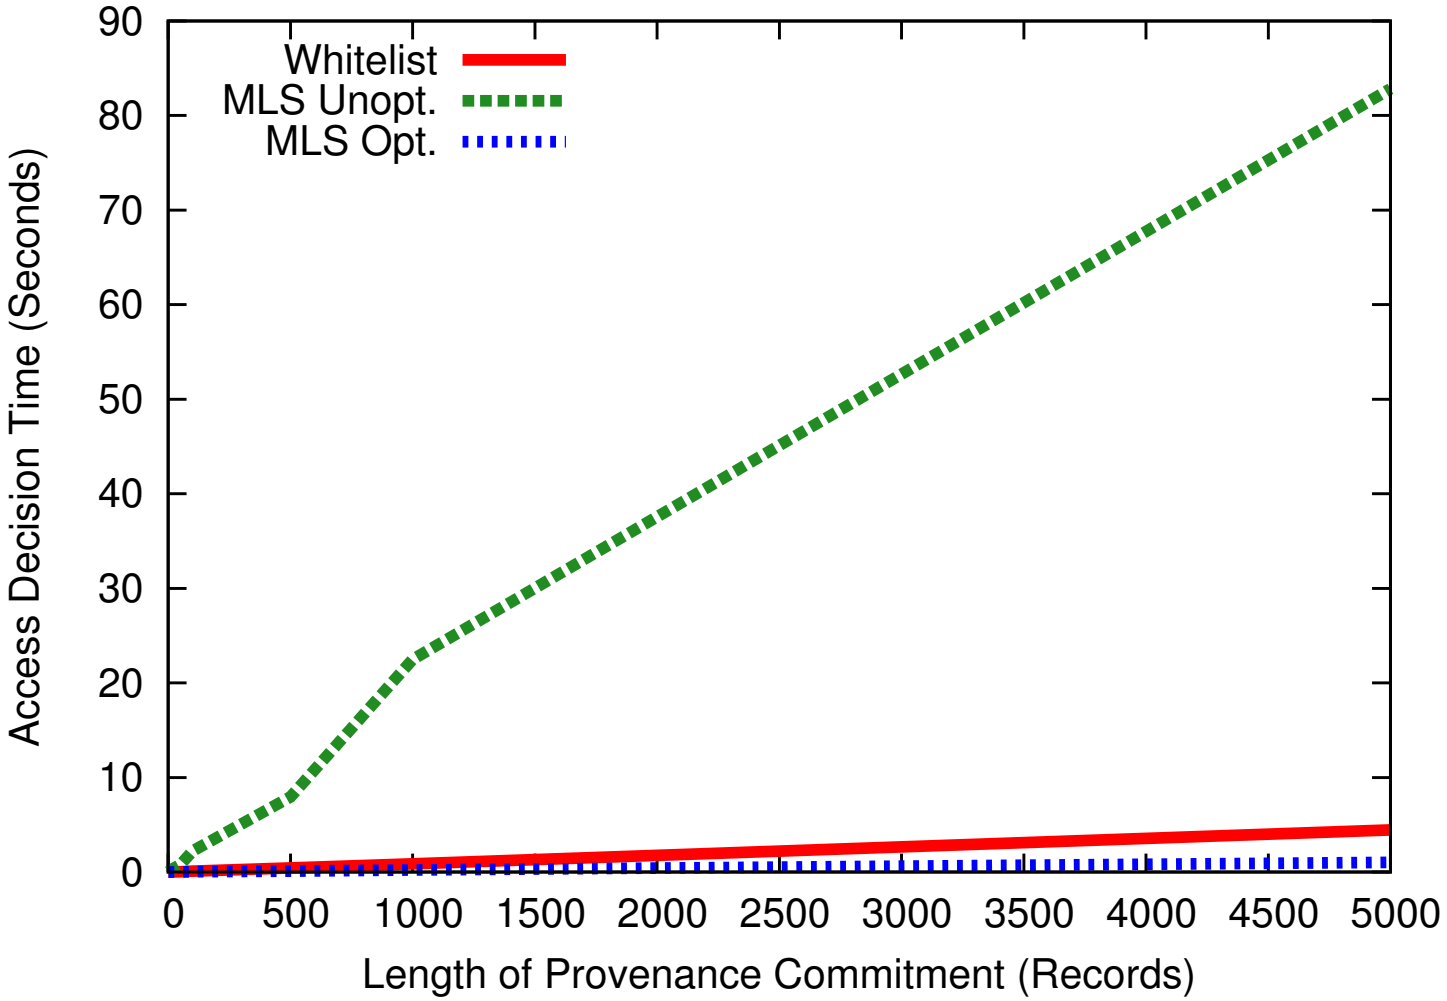

Since provenance is append-only, we cached previous access decisions in order to achieve amortized constant time.

Conclusion

UNIVERSITY OF OREGON

- *Cloud Provenance Authorities* bring us one step closer to secure, provenance-aware distributed applications.
- We don't need to wait on cloud providers to offer provenance services -- organizations can deploy *Cloud Provenance Authorities* using their own instances.
- Provenance applications such as access control can scale in the cloud when policy updates are infrequent.

Questions?

UNIVERSITY OF OREGON

Adam Bates

amb@cs.uoregon.edu

OSIRIS Oregon Systems Infrastructure Research & Information Security Laboratory