




Principles of Computer Systems: Naming

Prof. Katerina Argyraki

School of Computer & Communication Sciences

Introduction



A **name** is a way to refer to a **resource**

network arch. lab

0x12aadd

0x1348ad

0x12aadd

Zeinab Shmeis

PhD student

enrolled 2018

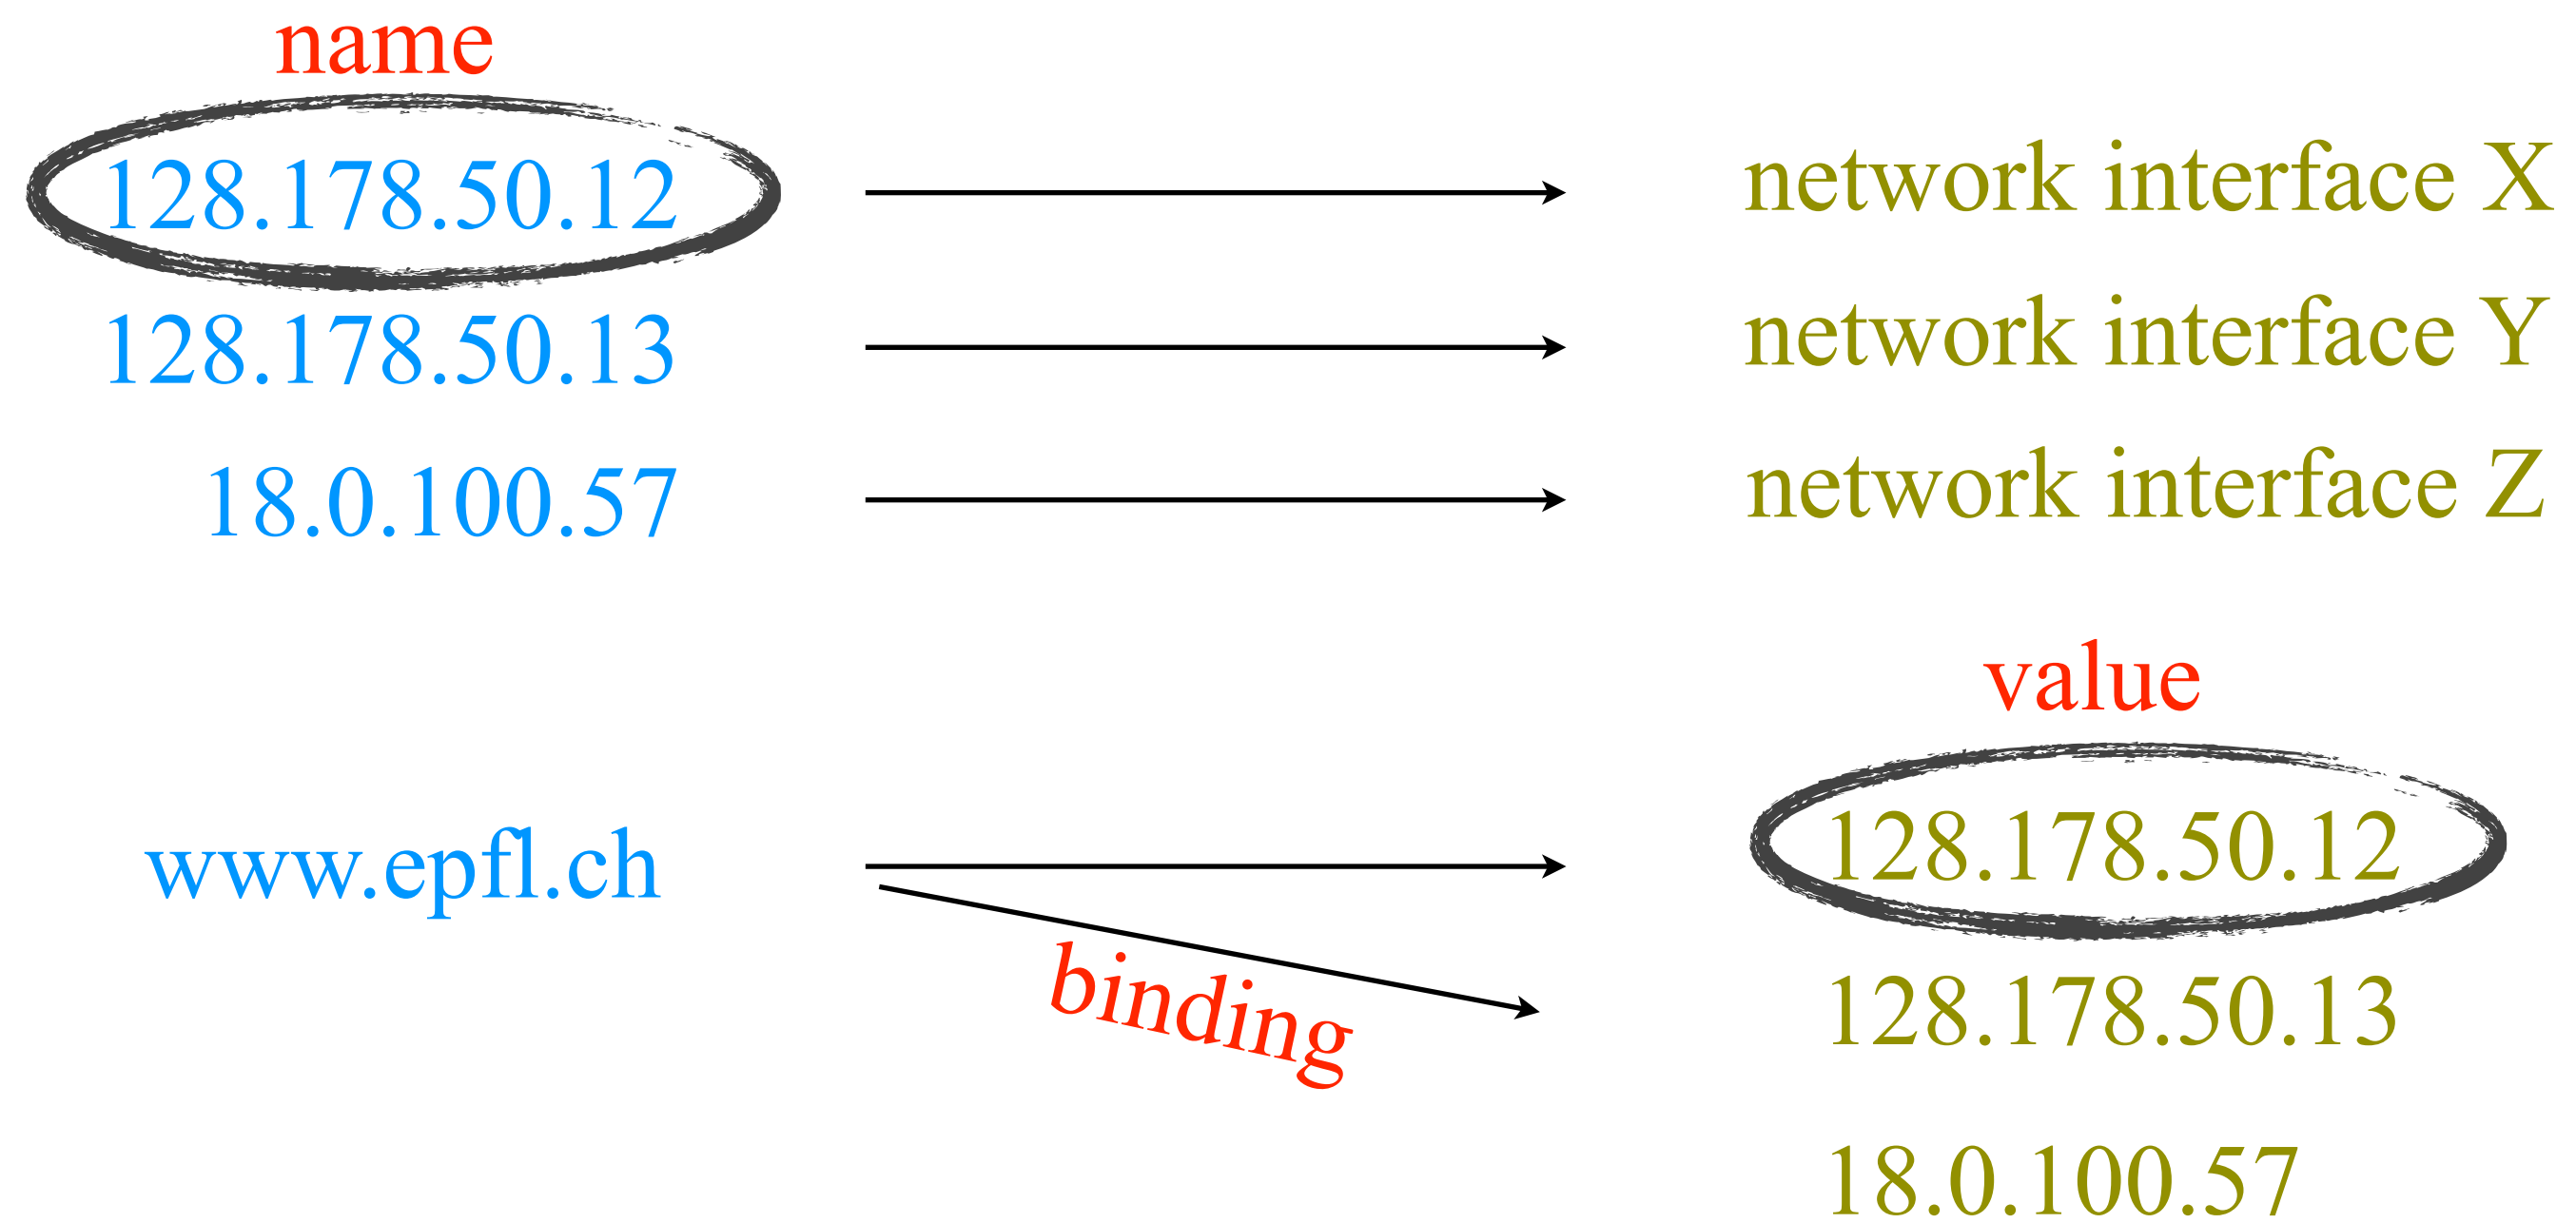
0x1348ad

Zhiyong Zhang

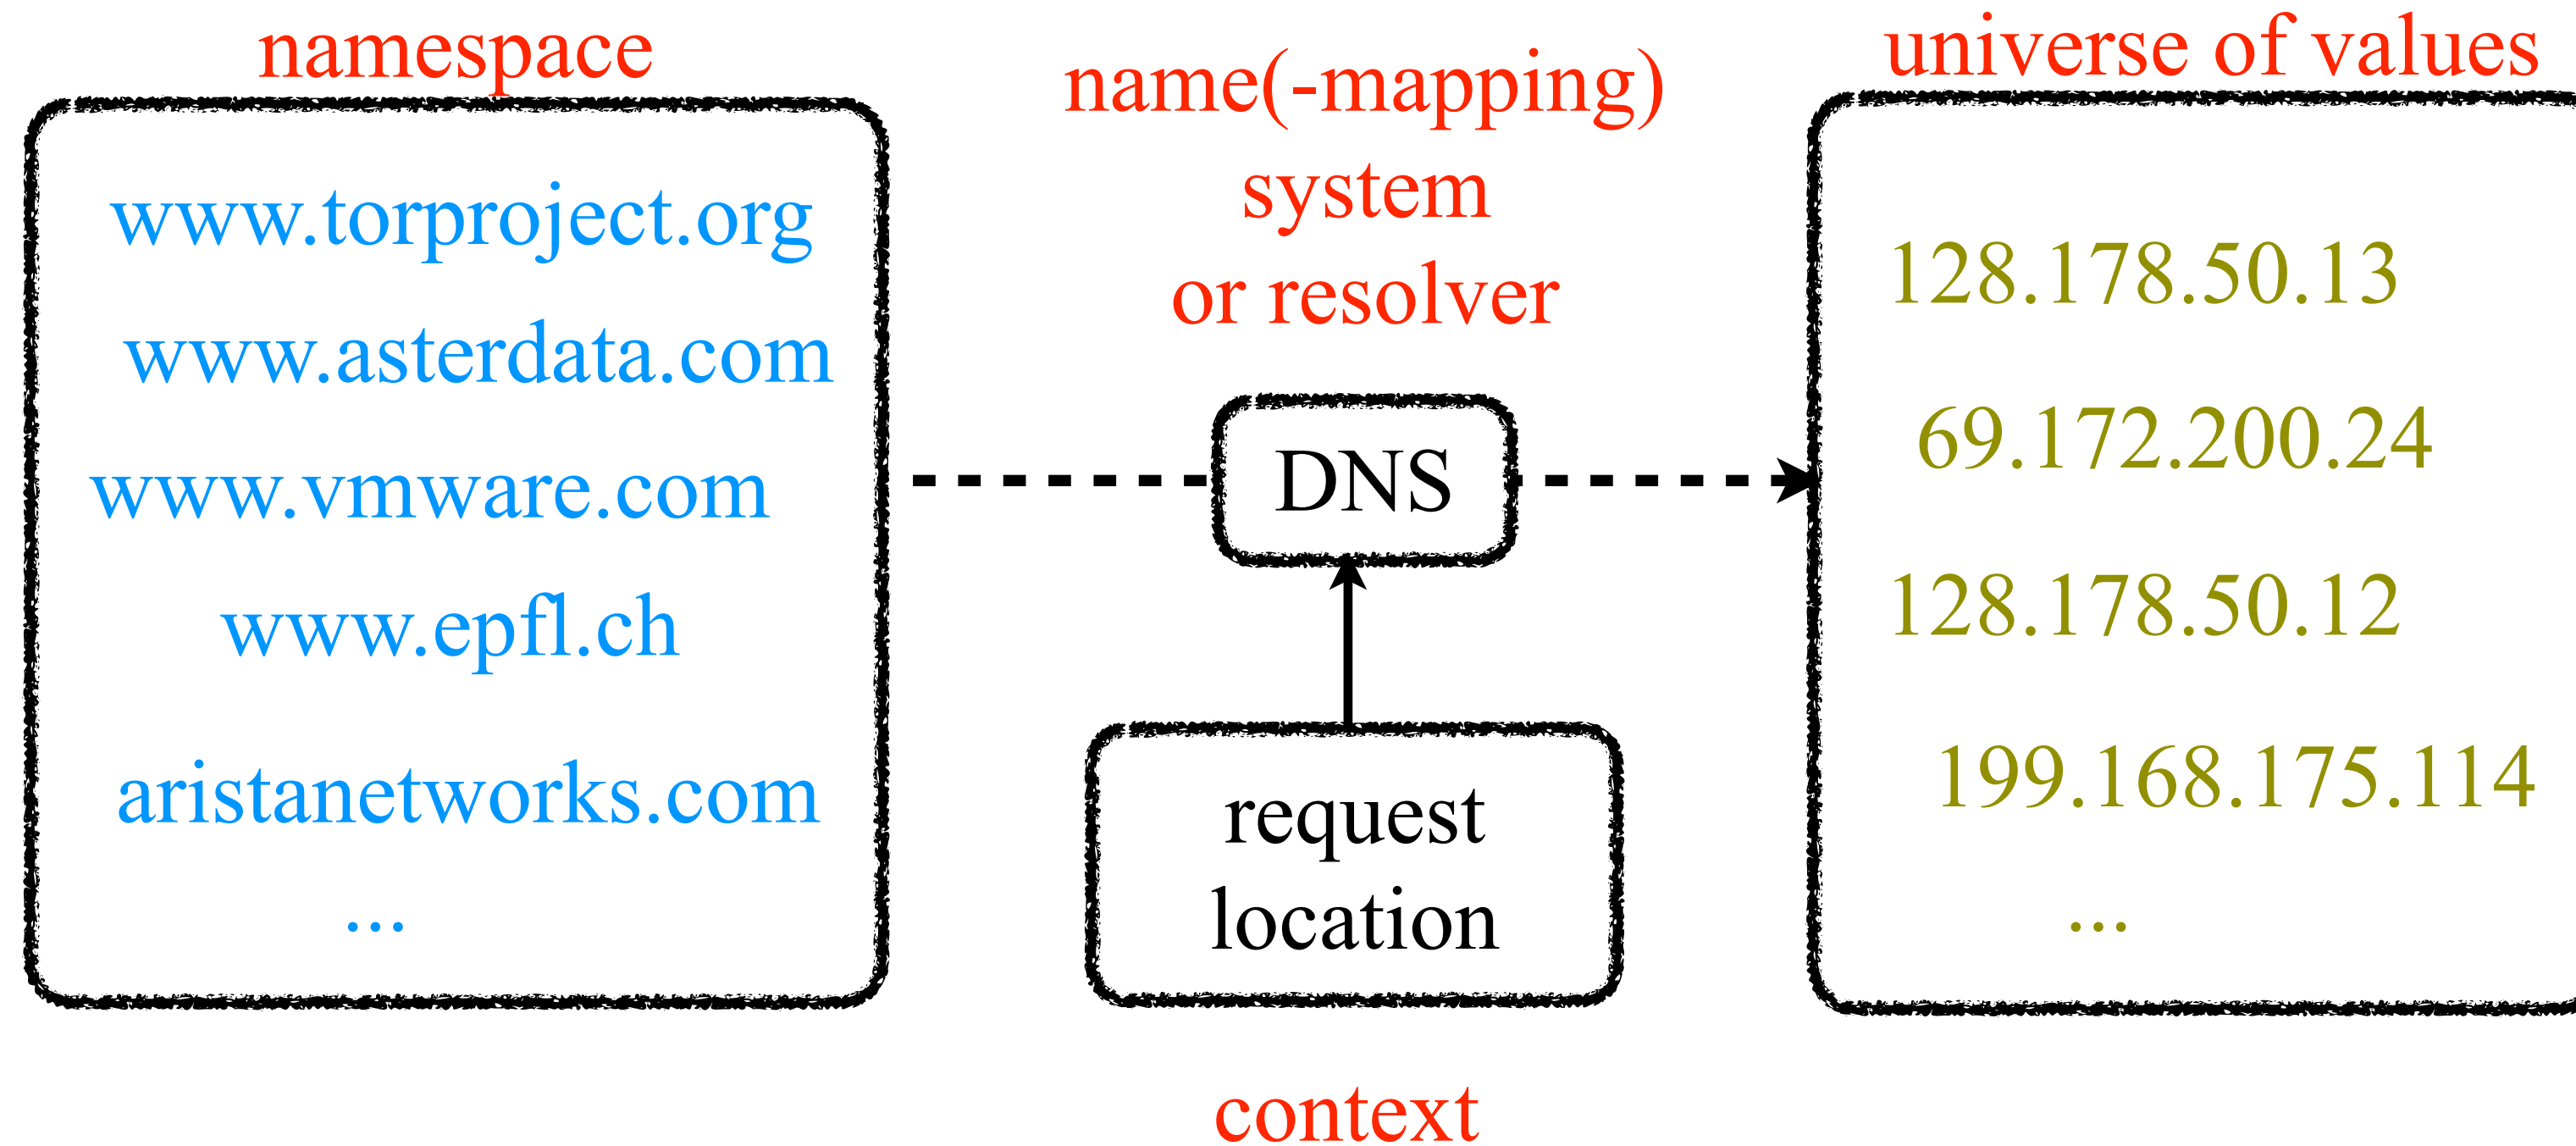
visiting student

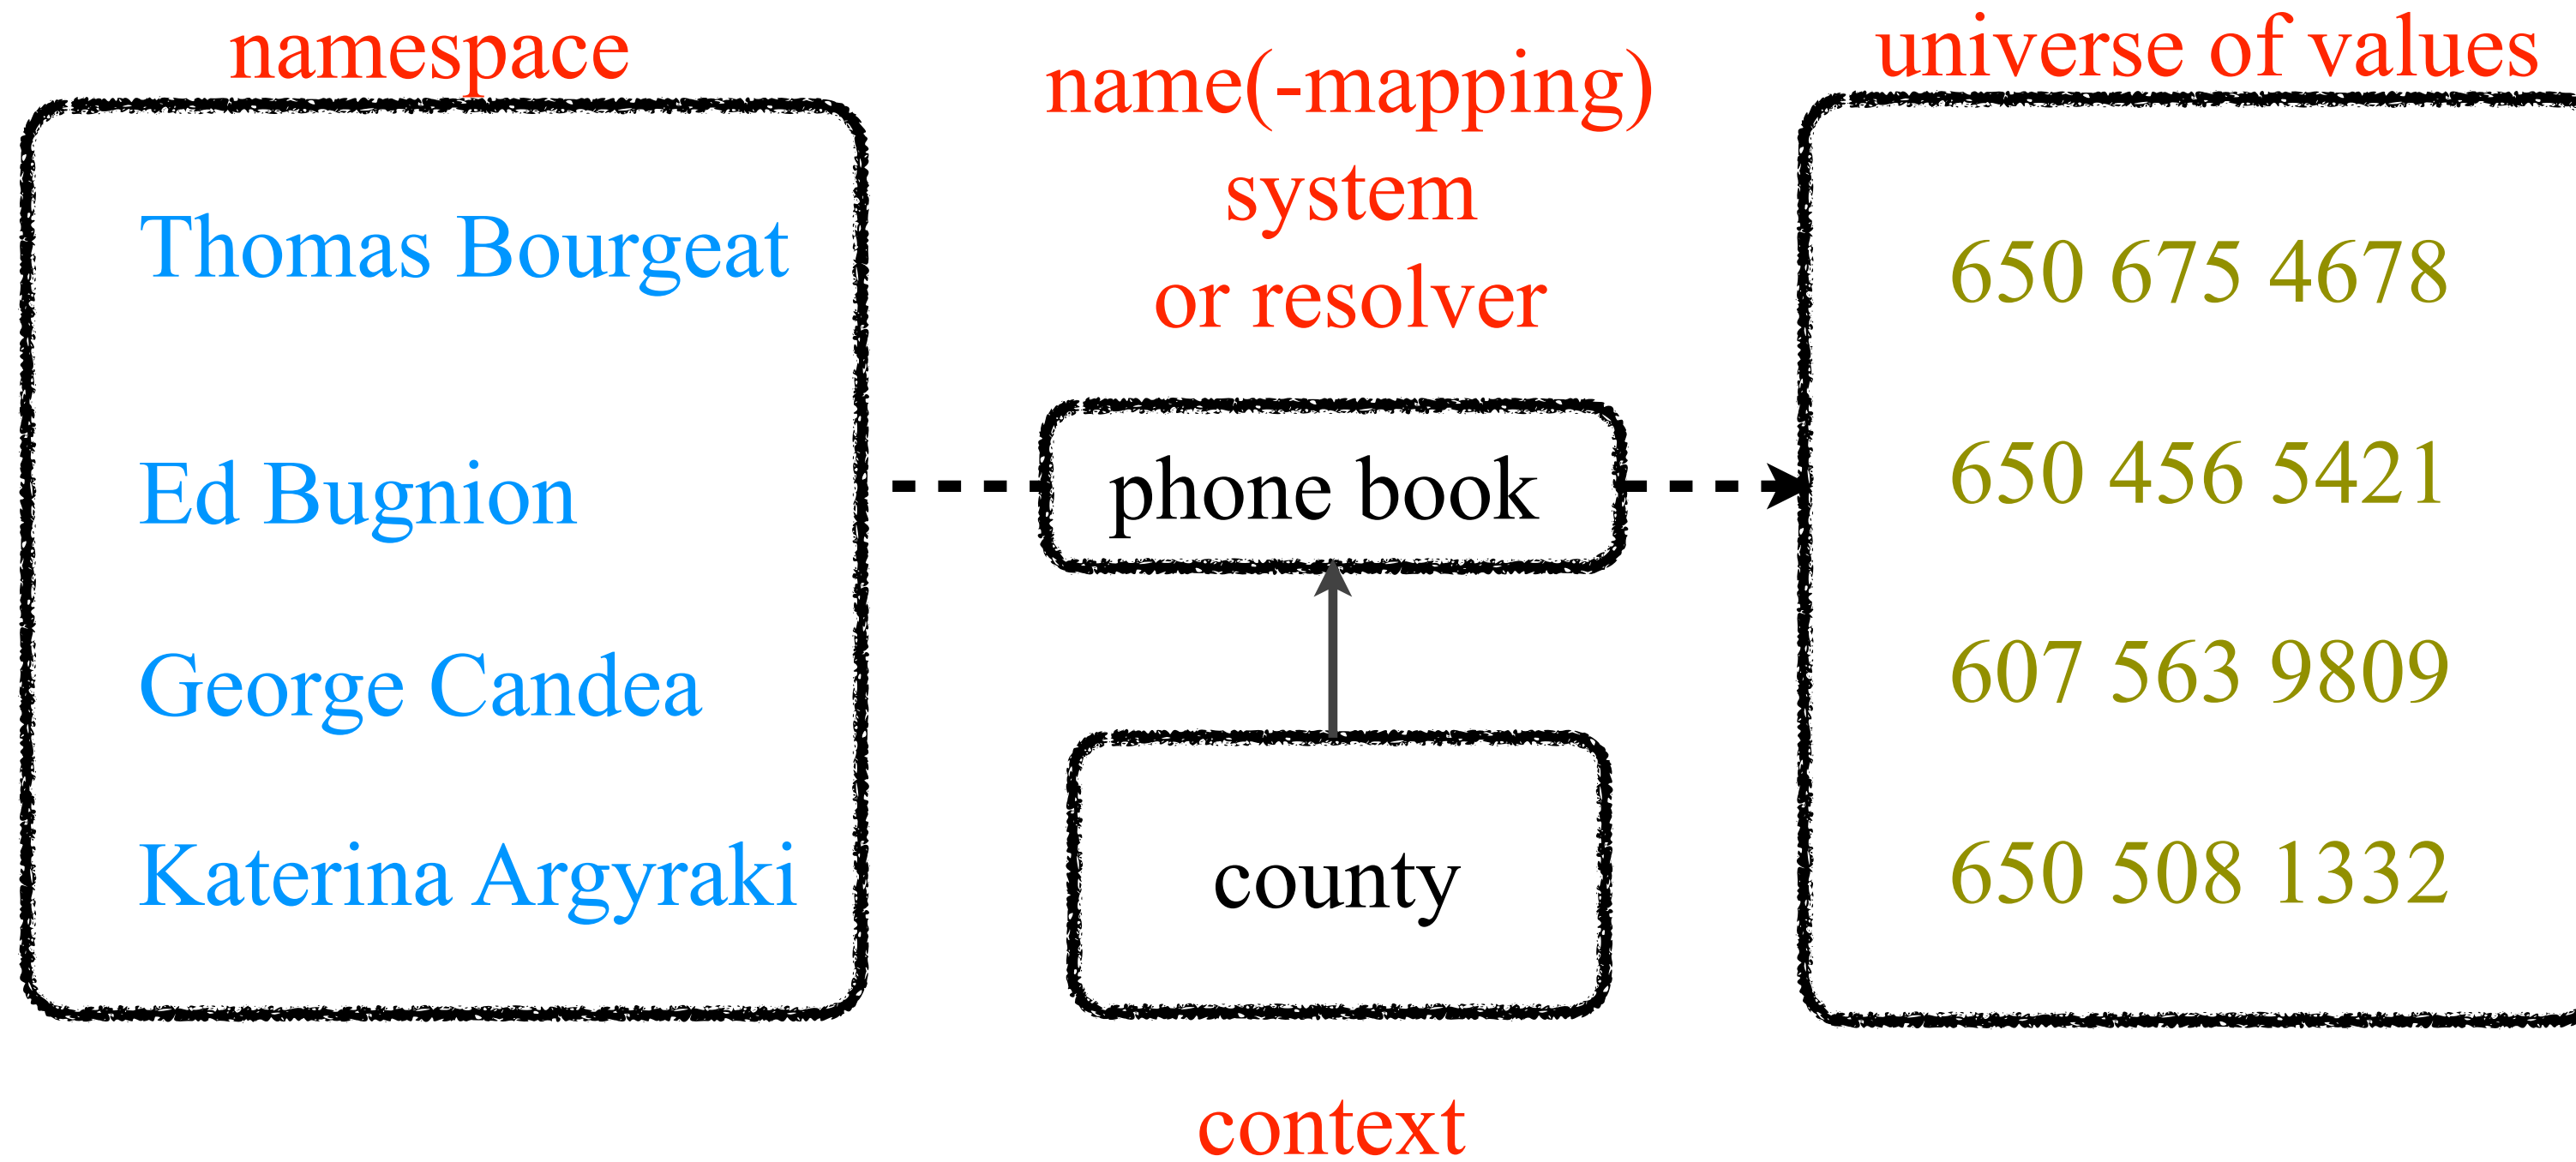
enrolled 2020

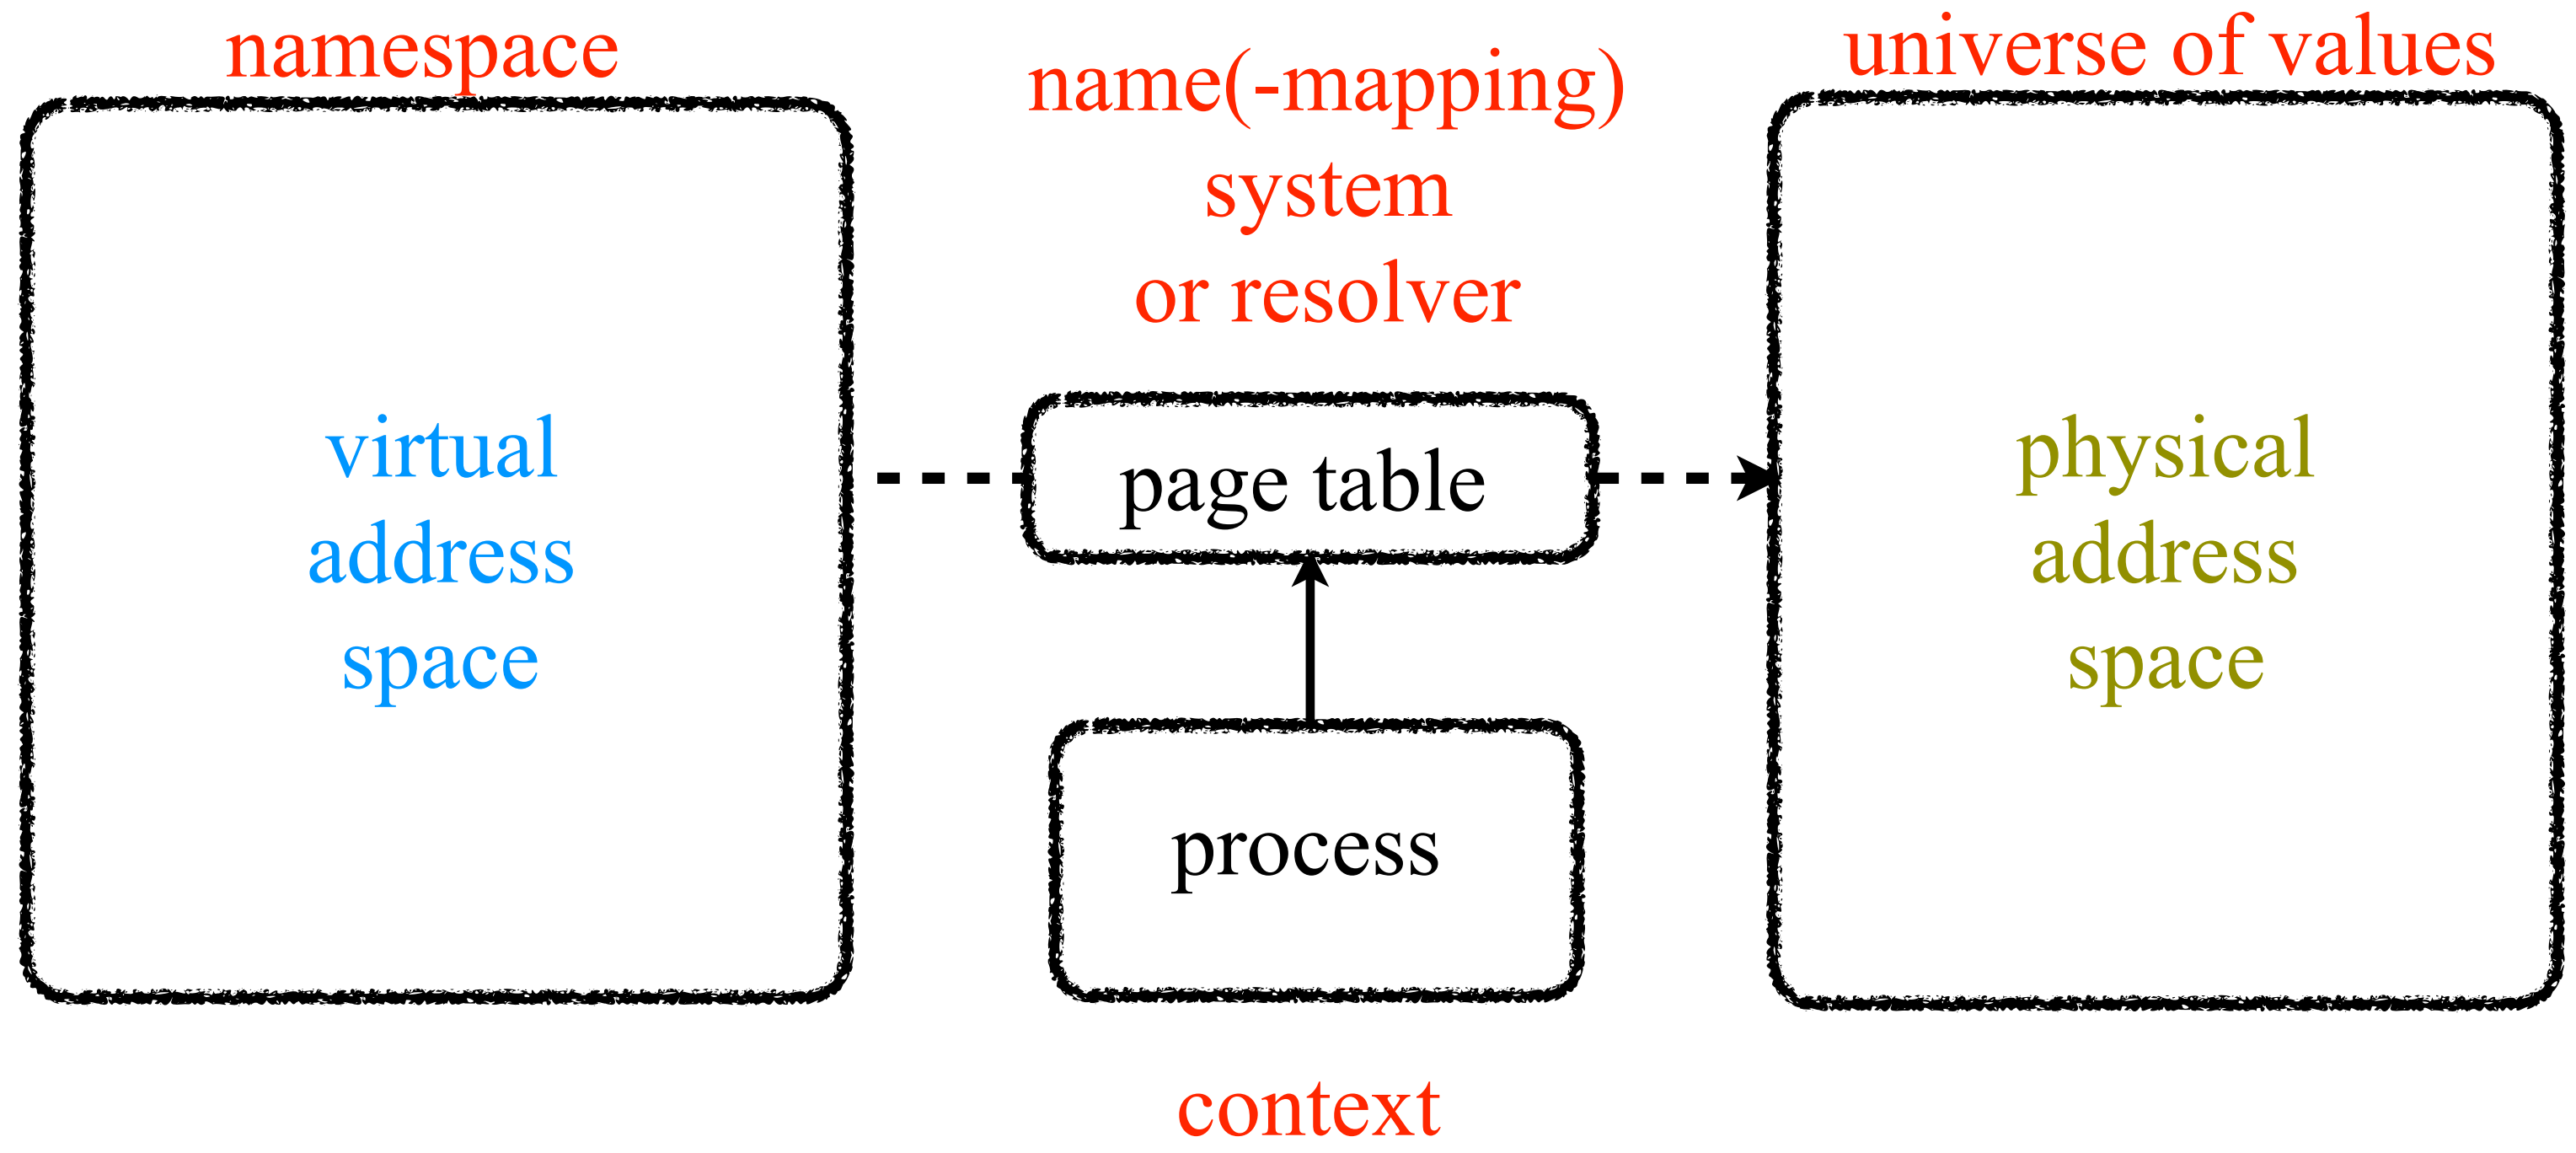
For efficient **communication** and **organization**



For **indirection**








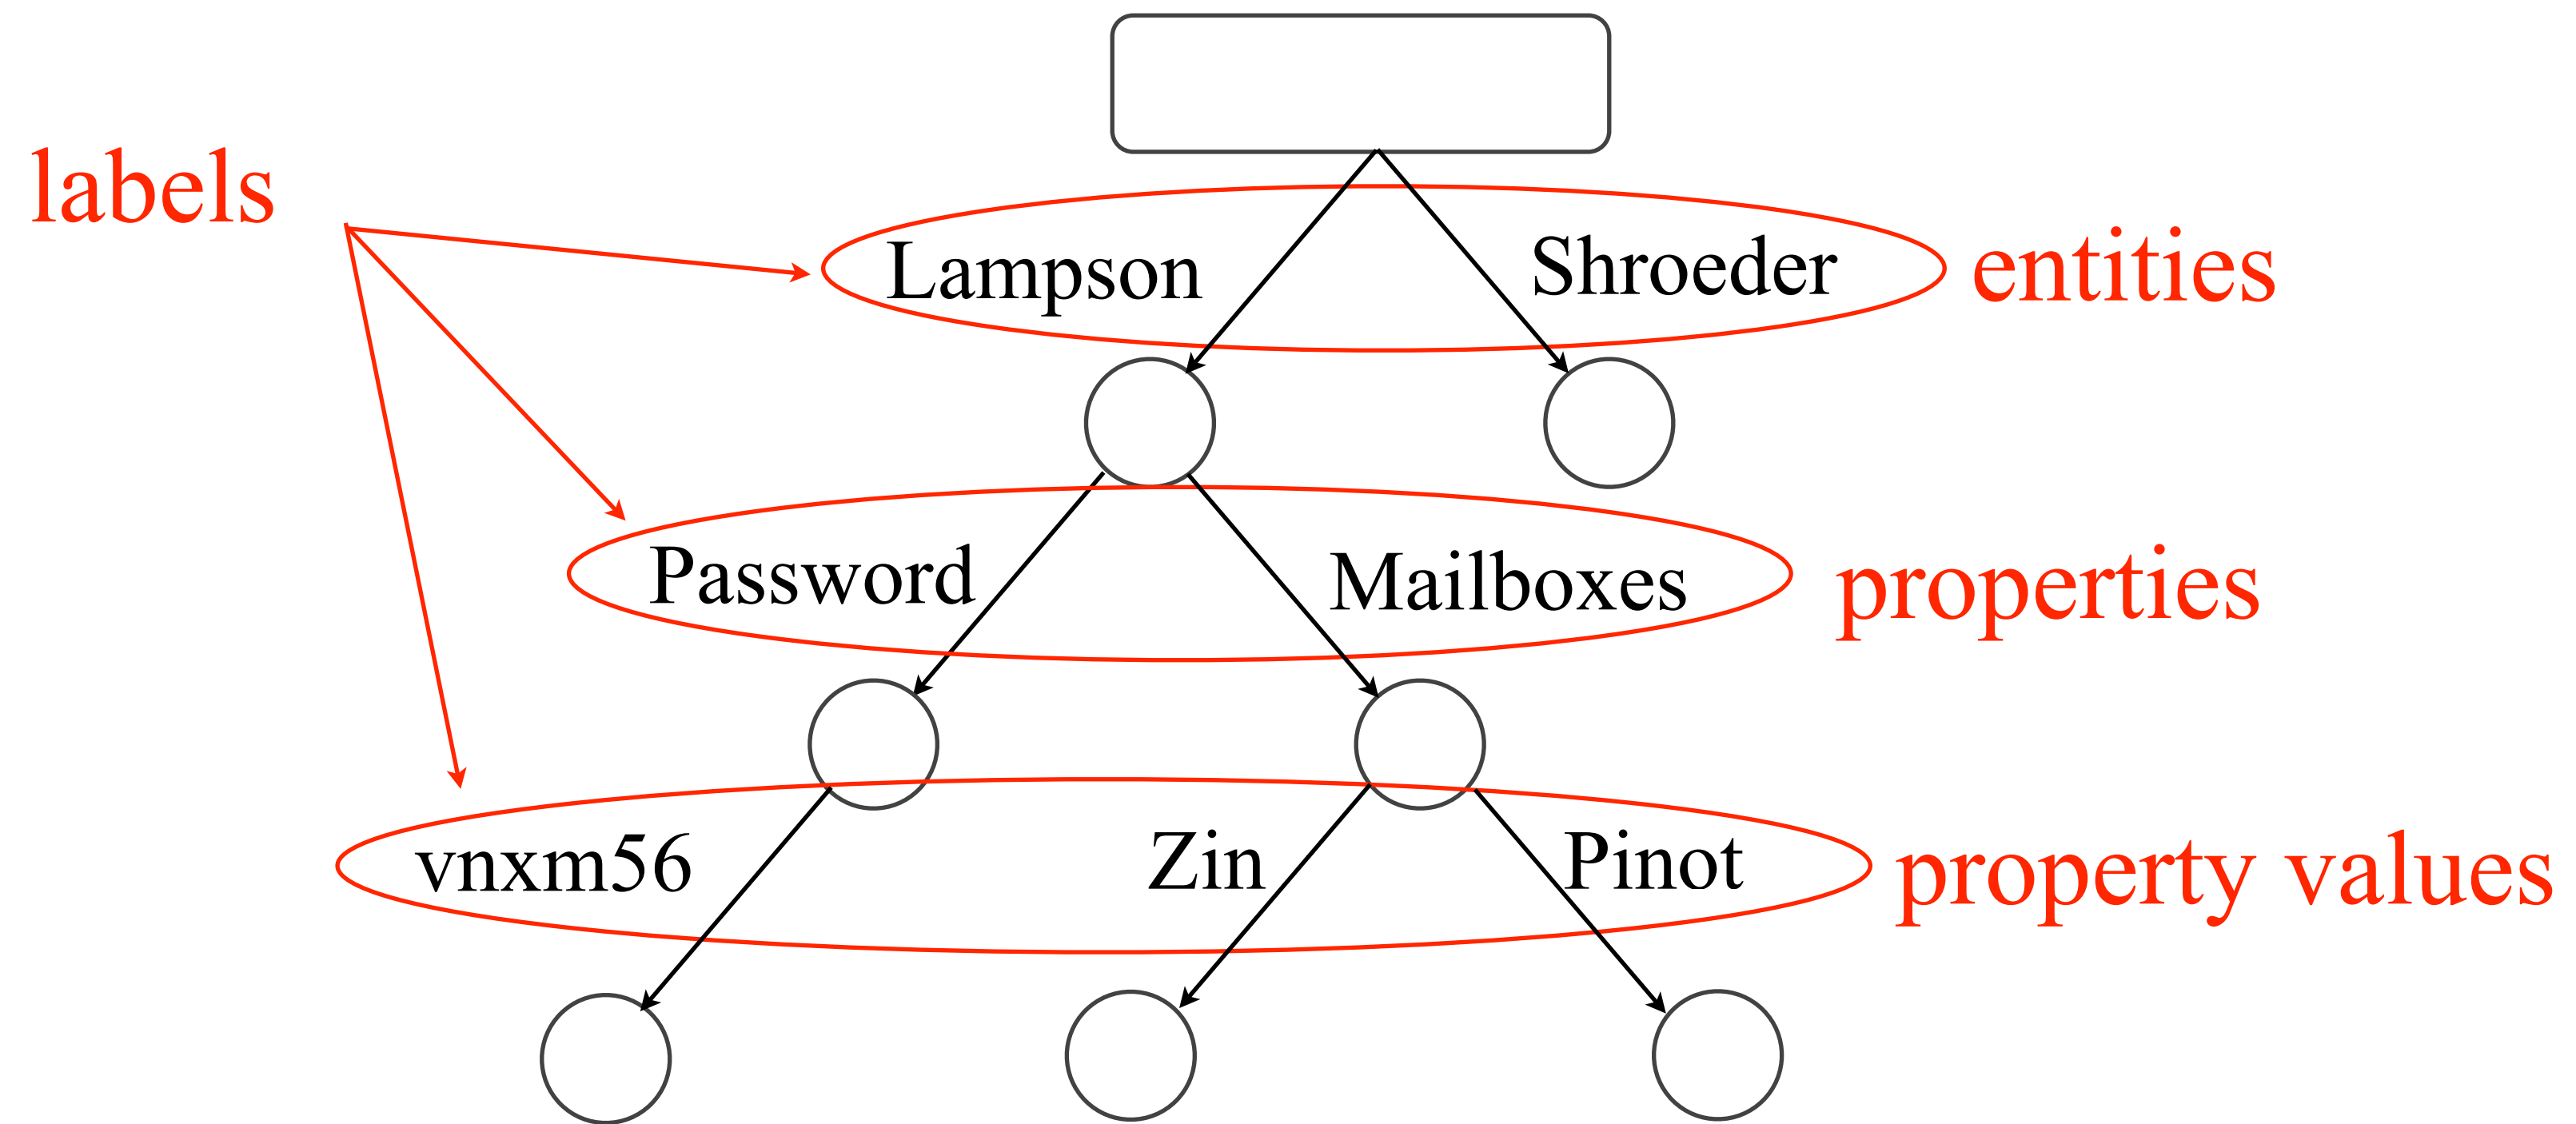
Name types

- **Private**: unique within a context
 - *e.g., a private IP address is unique within an organization*
- **Global**: unique across contexts
 - *e.g., a global IP address is unique within the Internet*
- **Hierarchical**: name relationship implies value relationship
 - *e.g., two IP addresses sharing the same prefix*
- **Flat**: name relationship implies nothing
 - *e.g., content IDs in Peer-to-Peer networks*

Designing a Global Name Service



- A name service maps **names** to **values**
- In this paper, a name is **hierarchical**, typically consisting of an entity and a property

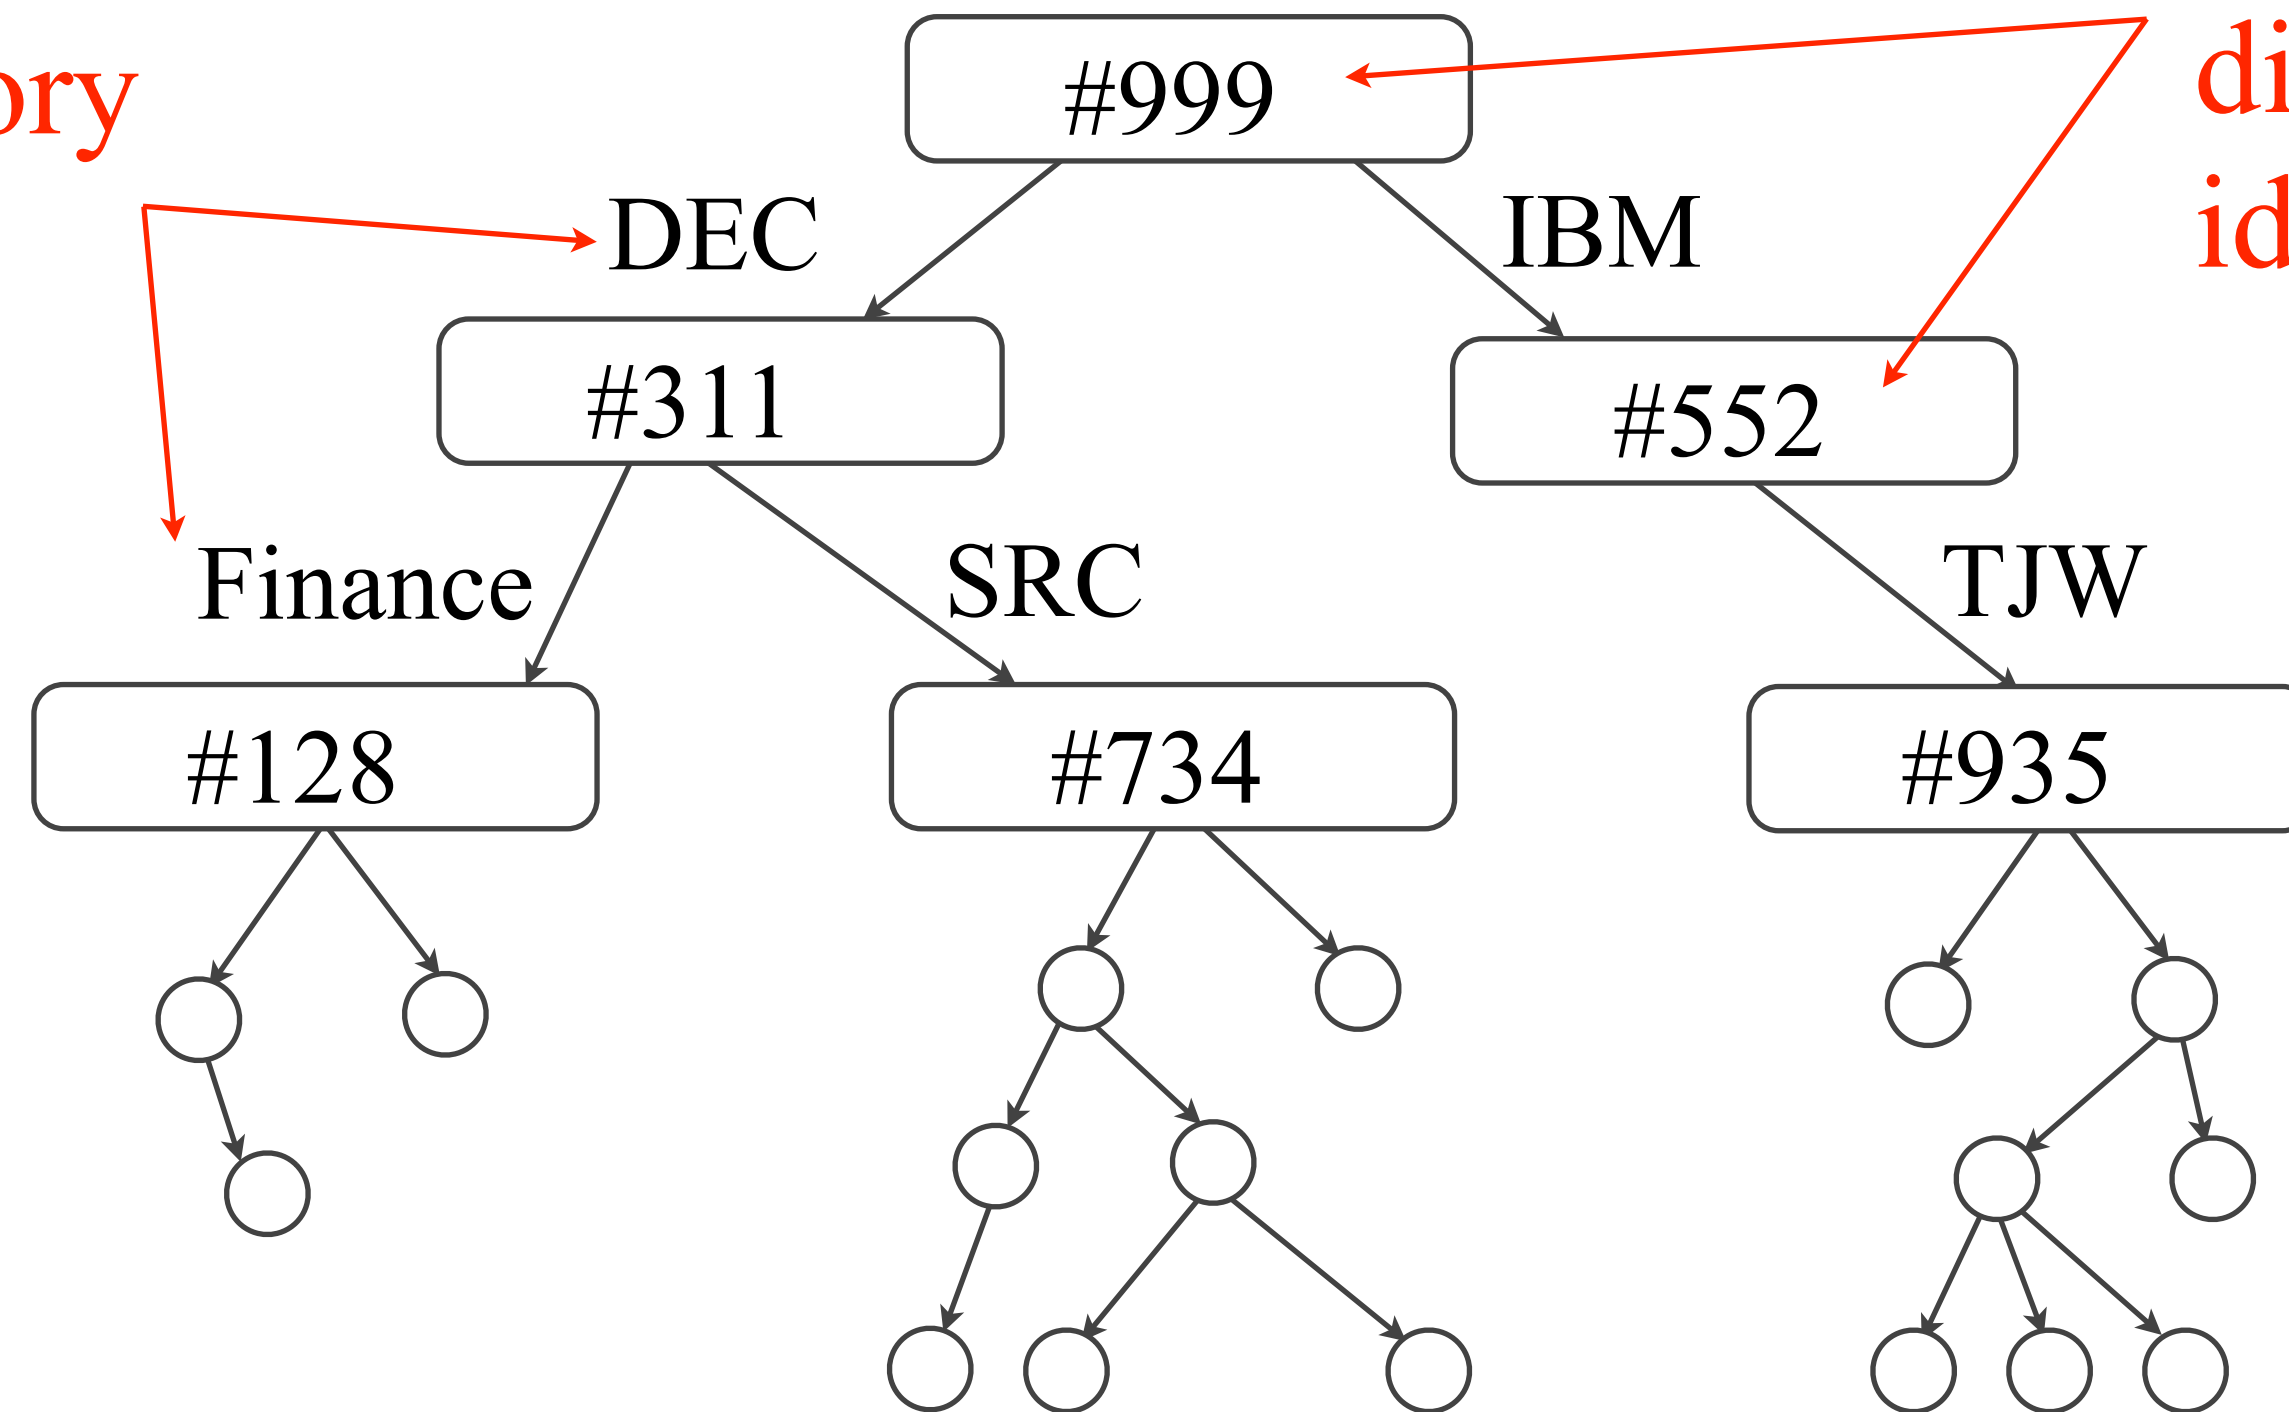


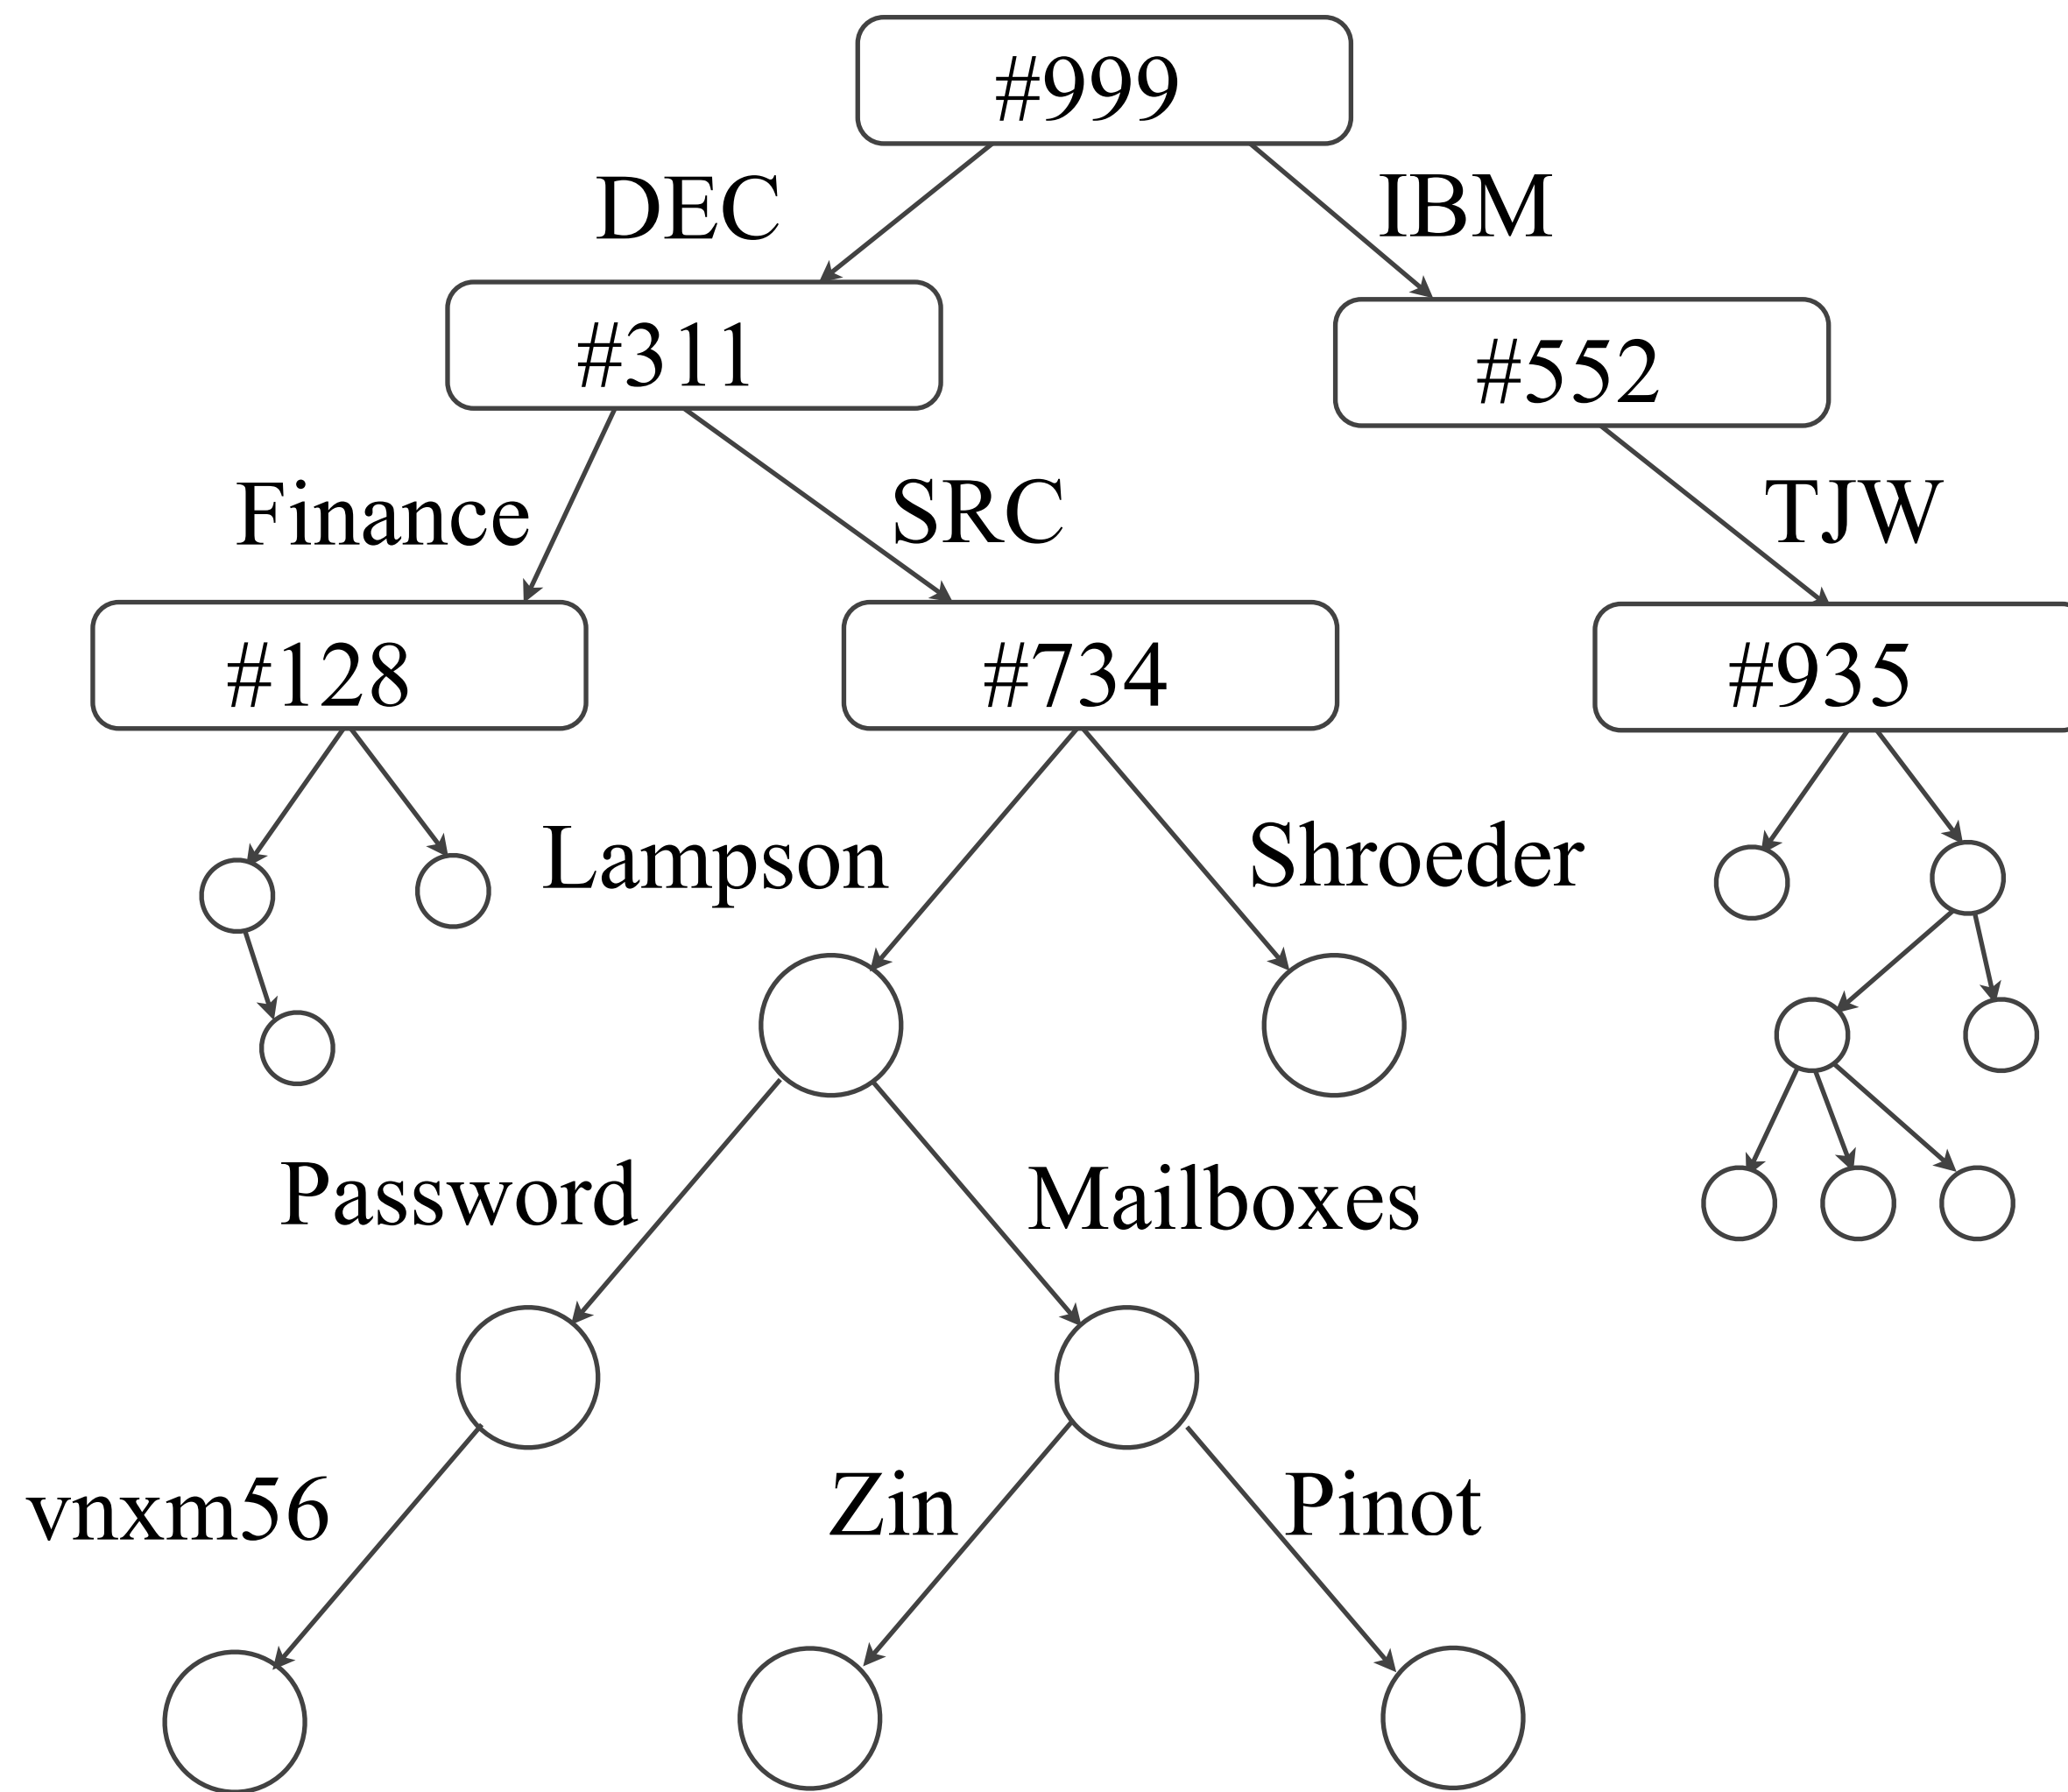
Design goal #1: scalability

- Must support an **arbitrary** number of names + administrative organizations

local
directory
names

global
directory
identifiers

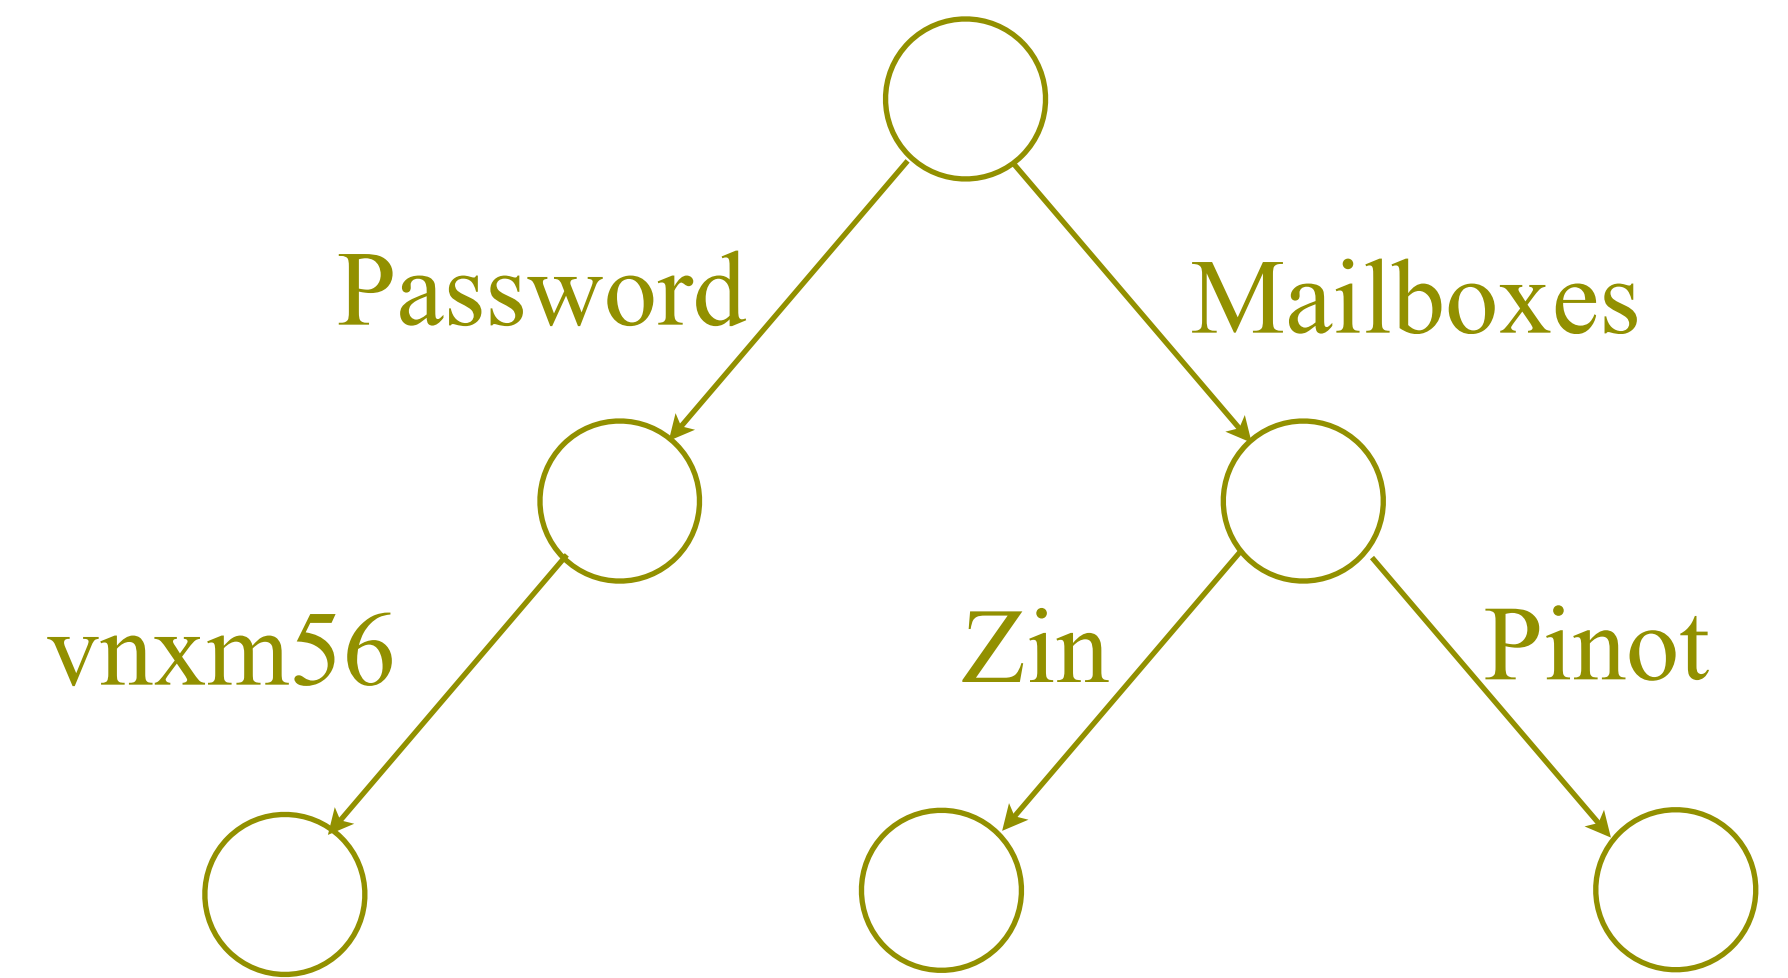




#999/DEC/SRC/Lampson/Password

vnxm56

#999/DEC/SRC/Lampson



Design goal #1: scalability


- Achieved through a **hierarchy** of directories, each potentially owned by a different entity, each with a **private namespace**

Design goal #2: **fault tolerance**

- The service should offer the same functionality even if N of its servers fail


servers

directory copies






servers


directory copies




timestamps




lastSweep: 2




lastSweep: 2



lastSweep: 2




lastSweep: 2




sweepTS: 5


lastSweep: 2




lastSweep: 2



lastSweep: 2




lastSweep: 2




sweepTS: 5


lastSweep: 2




lastSweep: 2



lastSweep: 2




lastSweep: 2




sweepTS: 5


lastSweep: 2




lastSweep: 2



lastSweep: 2




lastSweep: 2




sweepTS: 5


lastSweep: 2



lastSweep: 2




lastSweep: 25




sweepTS: 5


lastSweep: 2




lastSweep: 2



lastSweep: 2 5




lastSweep: 2




sweepTS: 5


lastSweep: 2



lastSweep: 25




lastSweep: 25




sweepTS: 5


lastSweep: 2



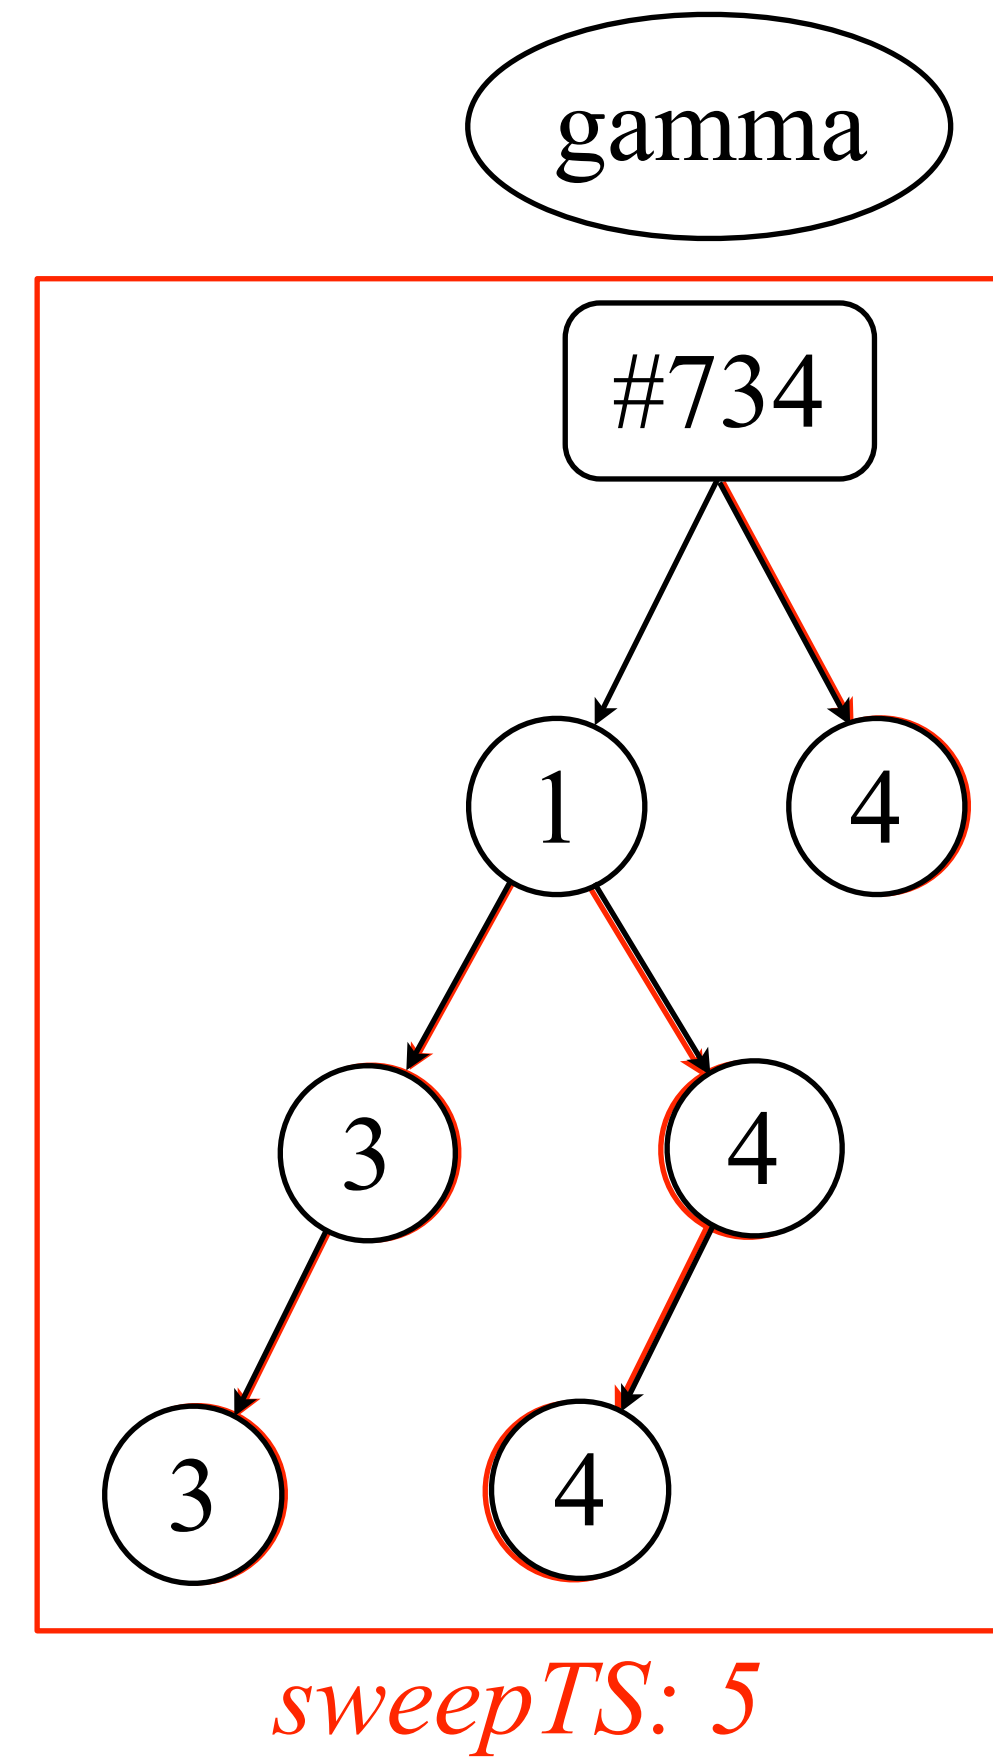
lastSweep: 25




lastSweep: 25




lastSweep: 2




lastSweep: 25




lastSweep: 25




lastSweep: 25




lastSweep: 25




lastSweep: 25




lastSweep: 25




lastSweep: 2




lastSweep: 2



lastSweep: 2




lastSweep: 2




sweepTS: 5


lastSweep: 2



lastSweep: 2




lastSweep: 2




sweepTS: 5


lastSweep: 2



lastSweep: 2




lastSweep: 2
nextTS: 5




sweepTS: 5


lastSweep: 2




lastSweep: 2



lastSweep: 2
nextTS: 5




lastSweep: 2




sweepTS: 5


lastSweep: 2



lastSweep: 2
nextTS: 5




lastSweep: 2
nextTS: 5




sweepTS: 5


lastSweep: 2




lastSweep: 2
nextTS: 5



lastSweep: 2
nextTS: 5




lastSweep: 2




sweepTS: 5


lastSweep: 2
nextTS: 5



lastSweep: 2
nextTS: 5




lastSweep: 2




sweepTS: 5


lastSweep: 2
nextTS: 5



lastSweep: 2
nextTS: 5




lastSweep: 2
nextTS: 5




sweepTS: 5


lastSweep: 2
nextTS: 5




lastSweep: 2
nextTS: 5



lastSweep: 2
nextTS: 5




lastSweep: 2
nextTS: 5




sweepTS: 5


lastSweep: 2
nextTS: 5



lastSweep: 2
nextTS: 5




lastSweep: 2 5
nextTS: 5




sweepTS: 5


lastSweep: 2
nextTS: 5




lastSweep: 2
nextTS: 5



lastSweep: 2 5
nextTS: 5




lastSweep: 2
nextTS: 5




sweepTS: 5


lastSweep: 2
nextTS: 5



lastSweep: 25
nextTS: 5




lastSweep: 25
nextTS: 5




sweepTS: 5


lastSweep: 2
nextTS: 5




lastSweep: 2 5
nextTS: 5




lastSweep: 2 5
nextTS: 5




lastSweep: 2
nextTS: 5




lastSweep: 25
nextTS: 5




lastSweep: 25
nextTS: 5




lastSweep: 25
nextTS: 5




lastSweep: 25
nextTS: 5



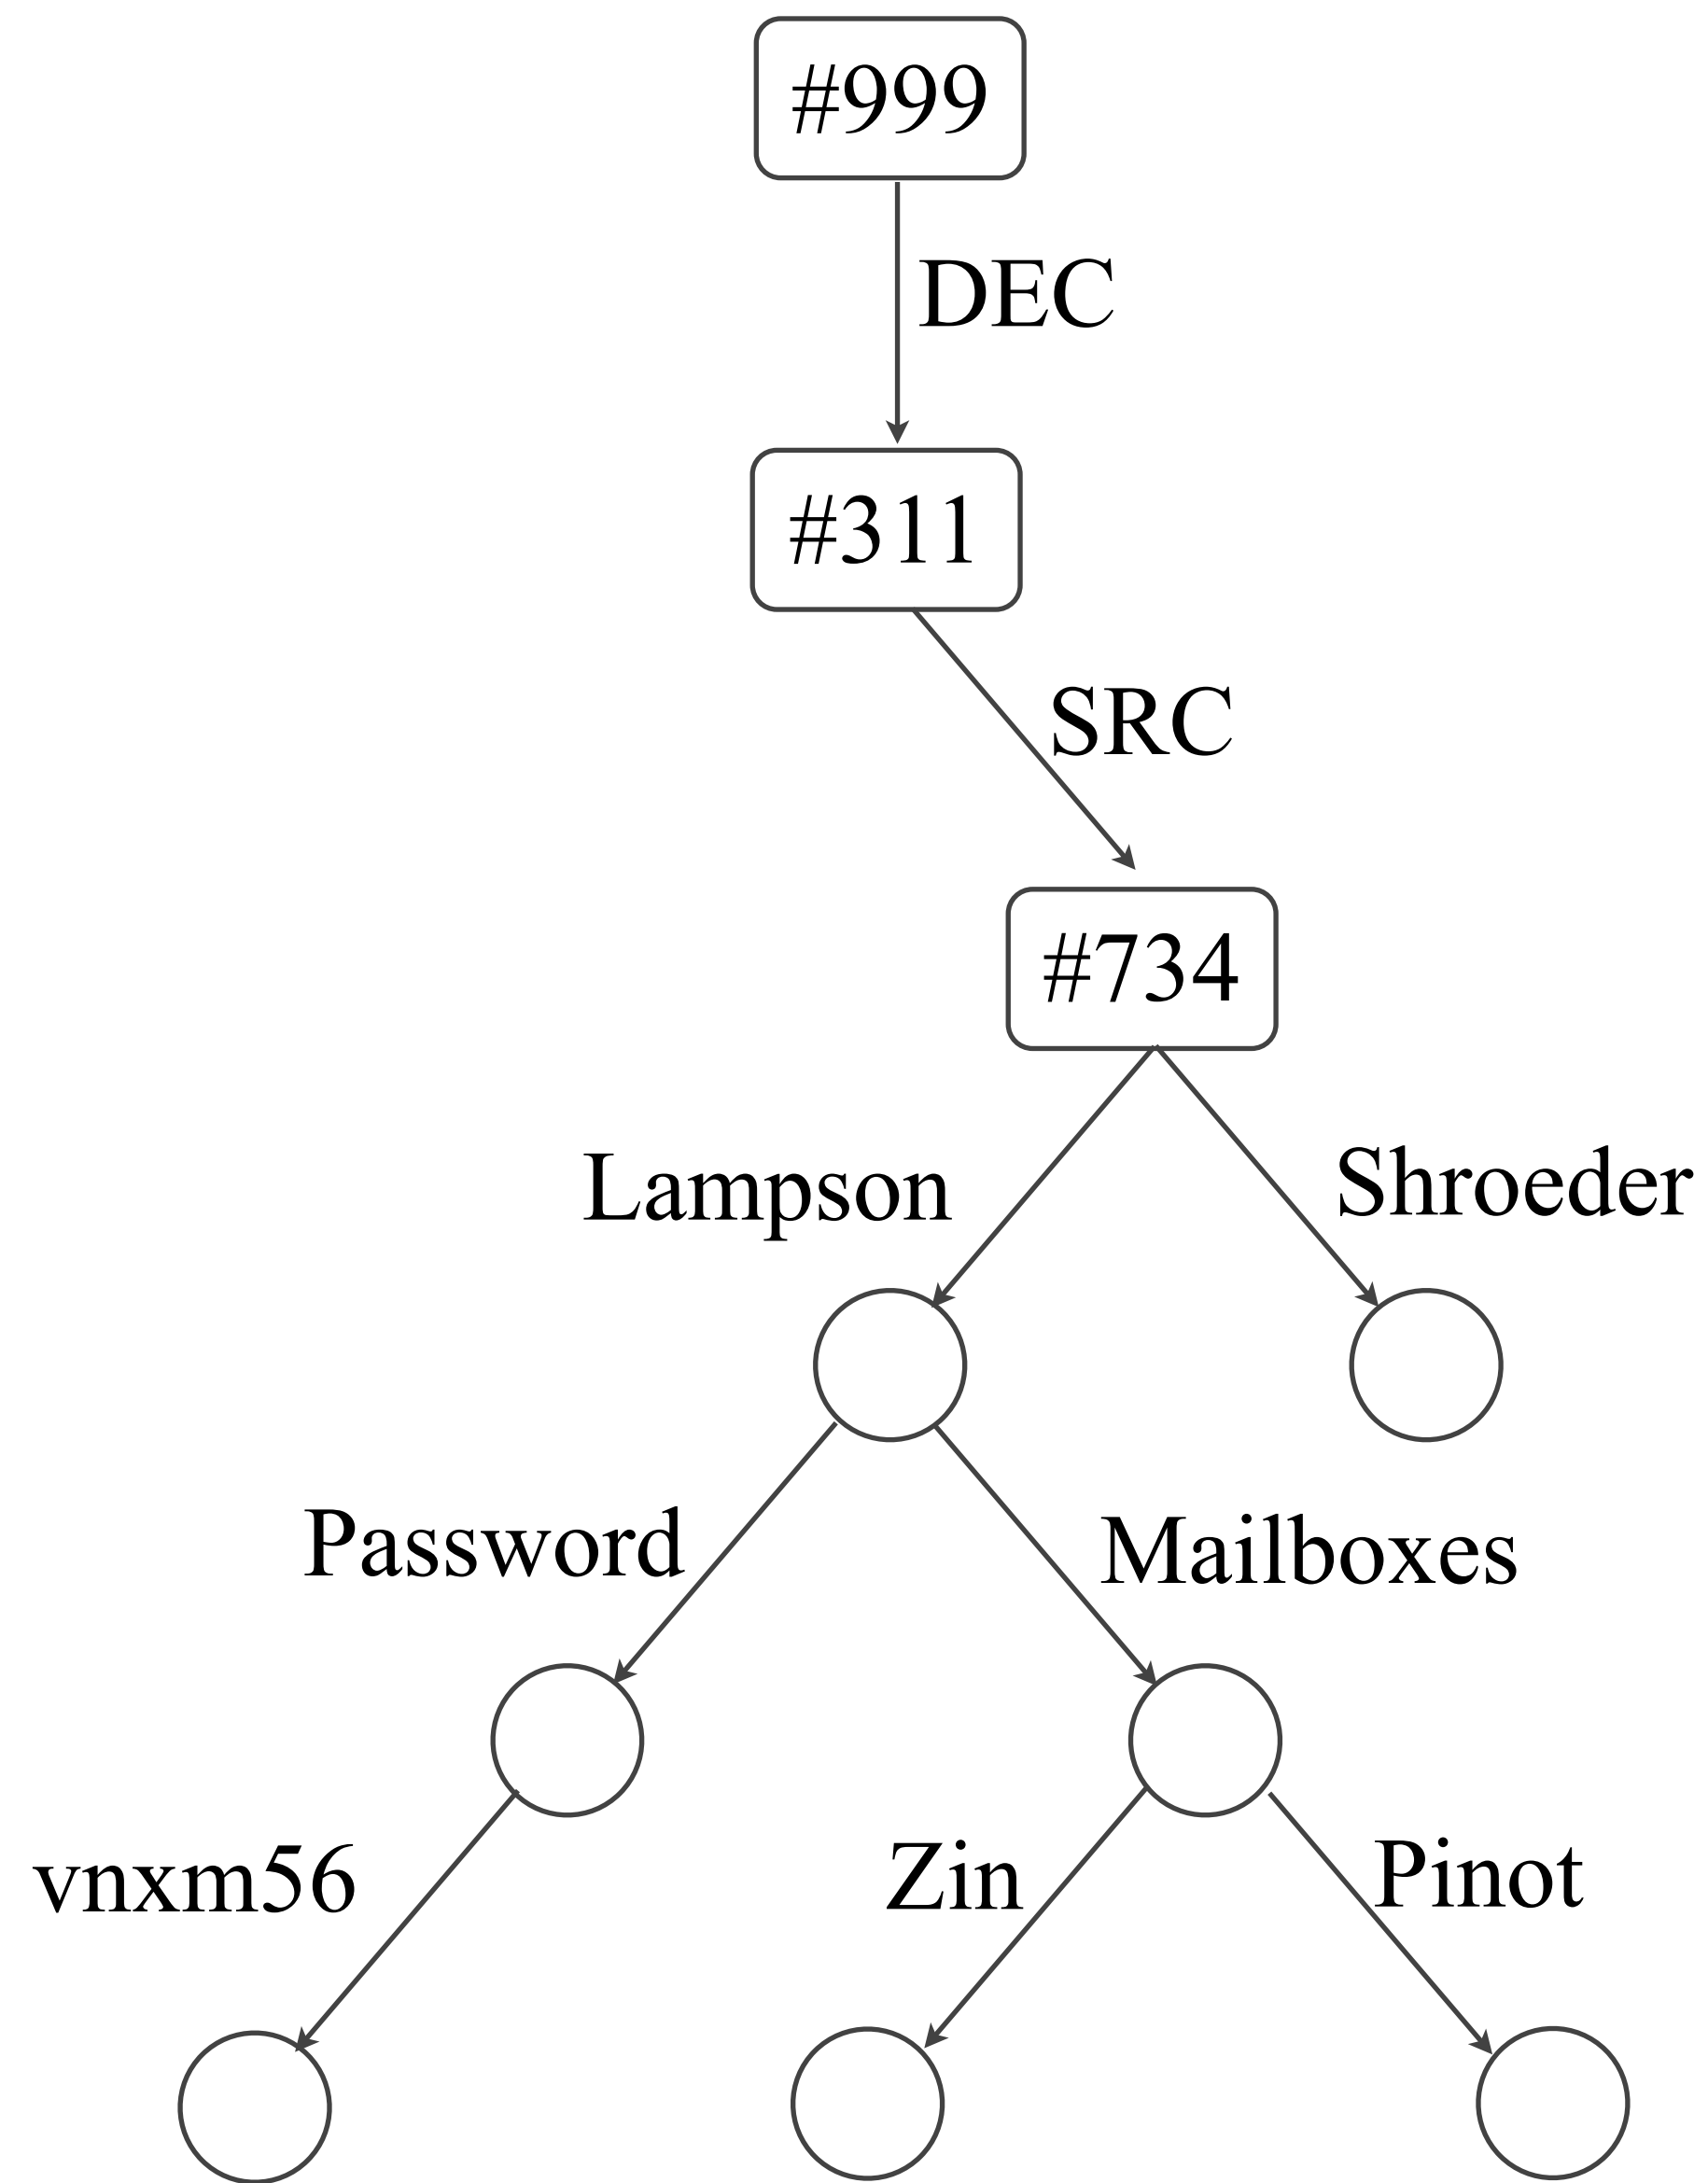
lastSweep: 25
nextTS: 5



lastSweep: 25
nextTS: 5



Name lookup

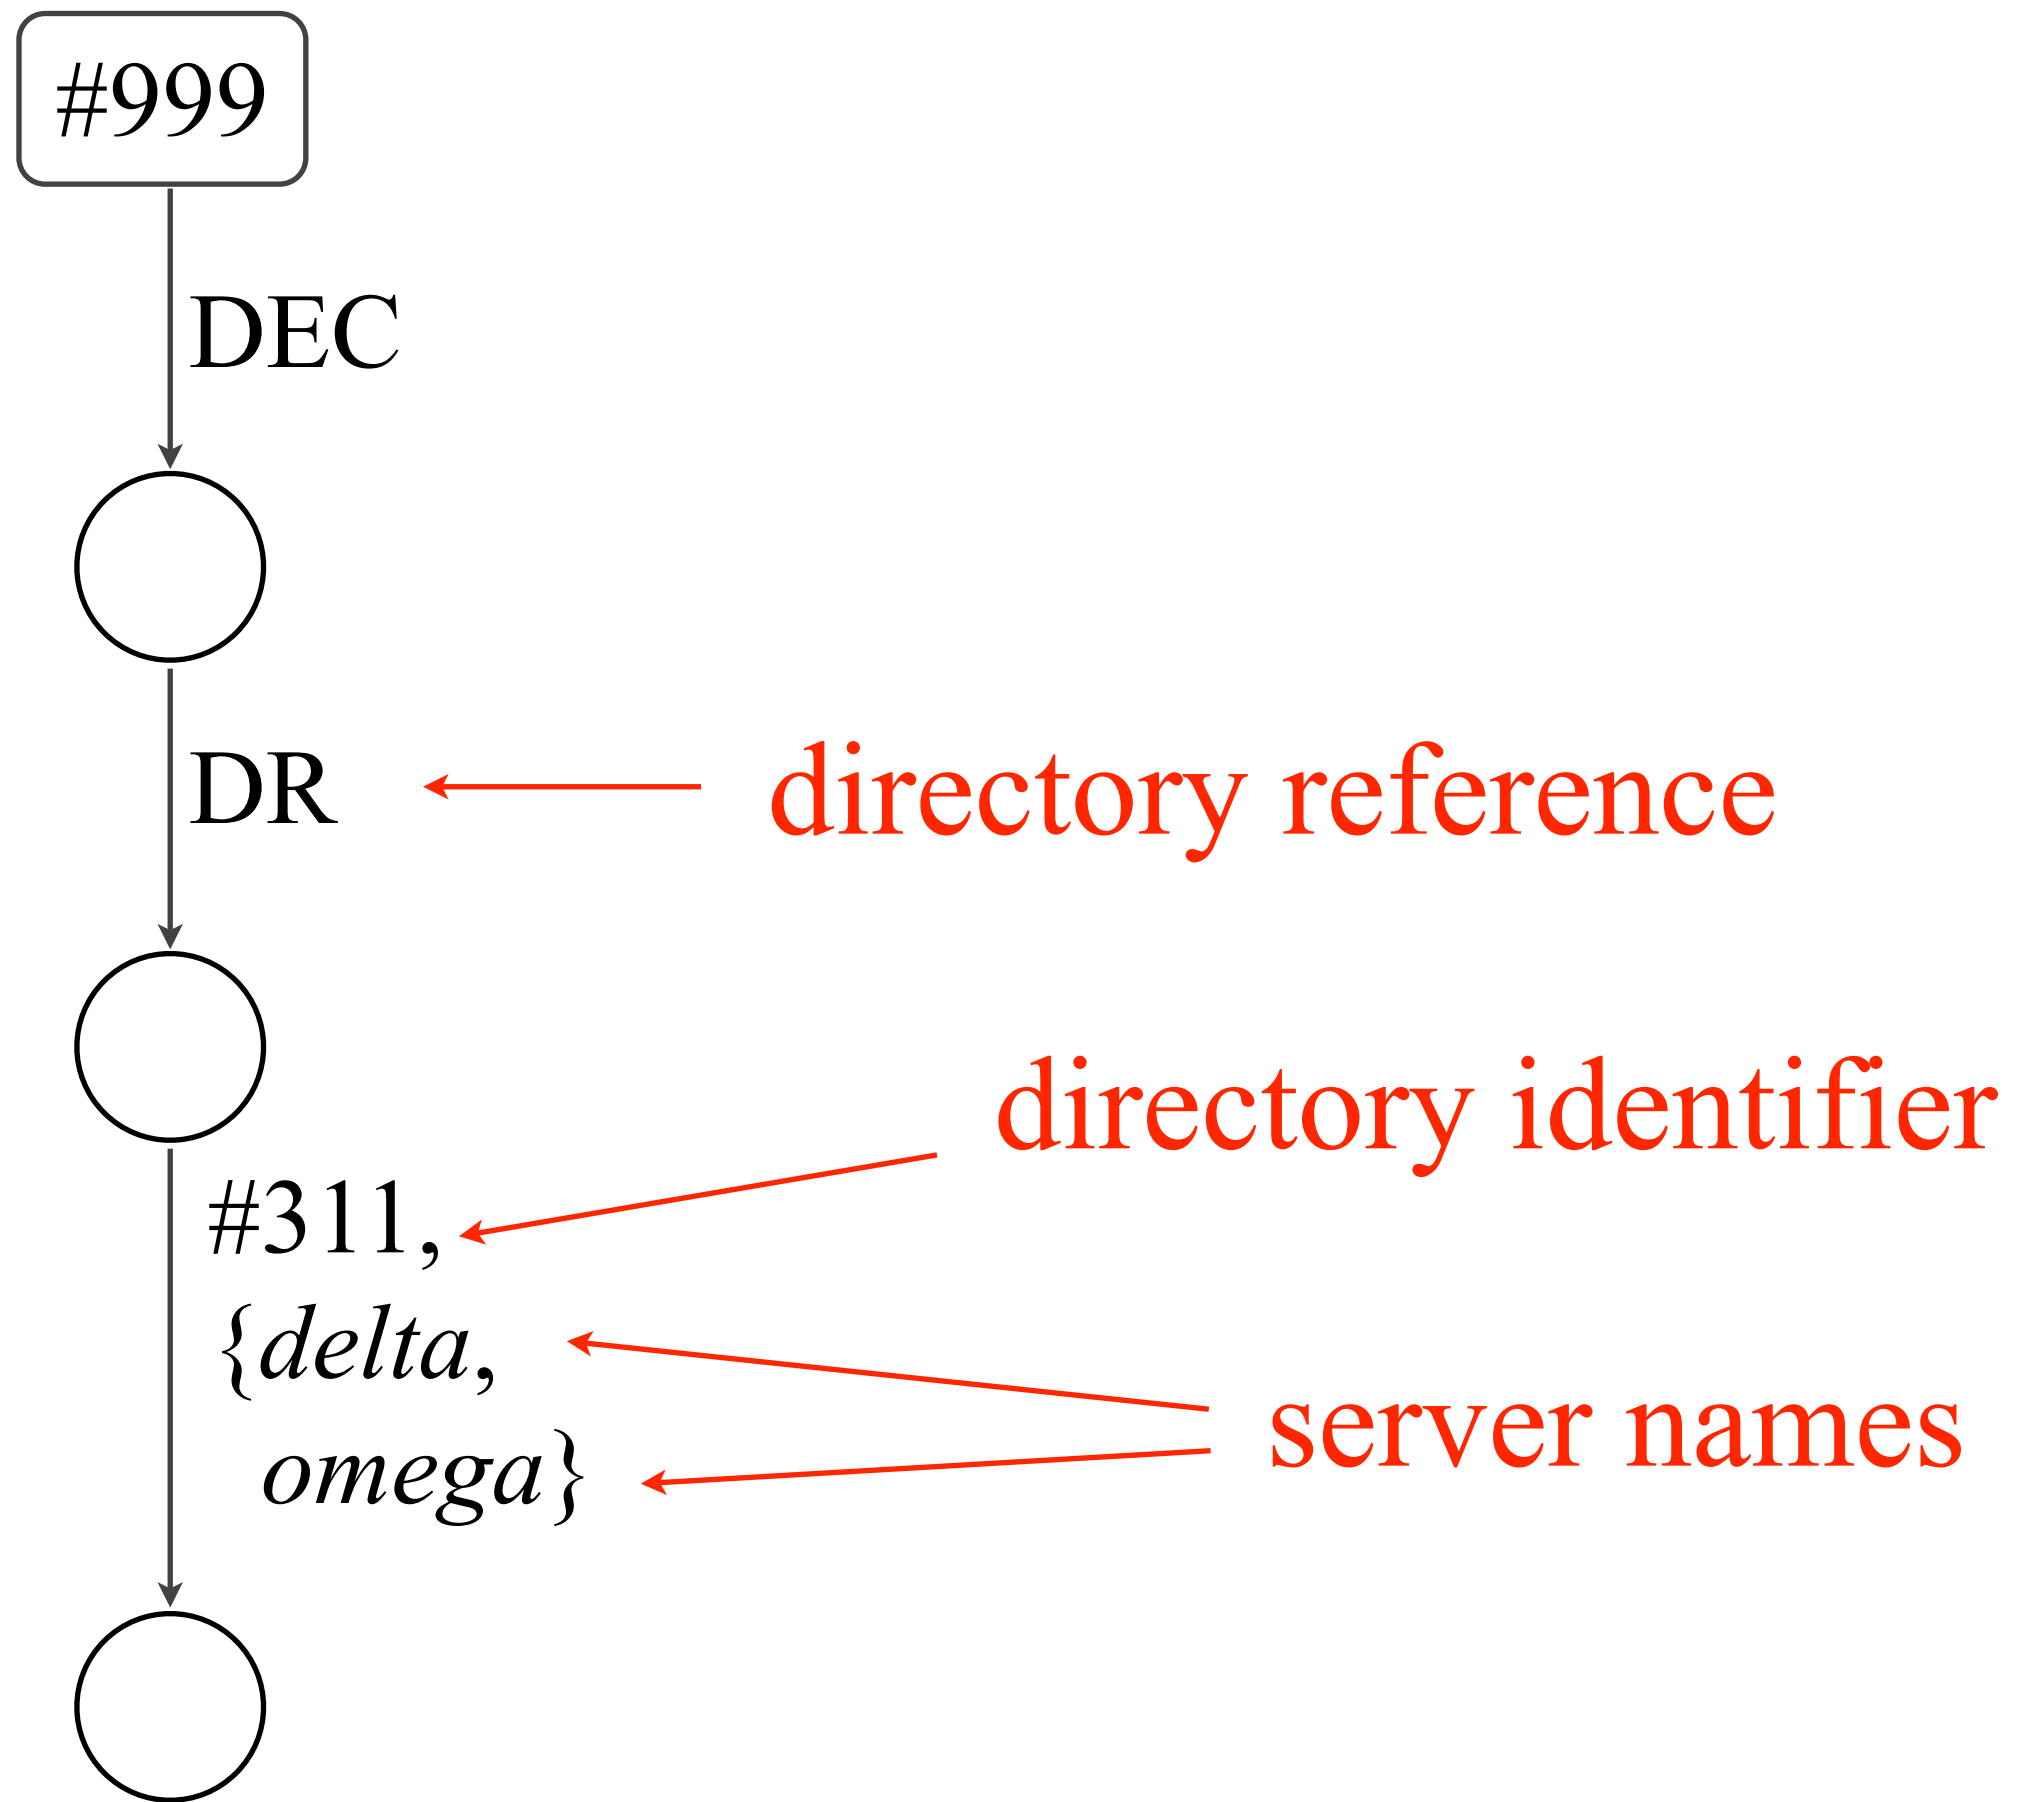


Name

Value

#999/DEC/SRC/Lampson/Password

vnxm56



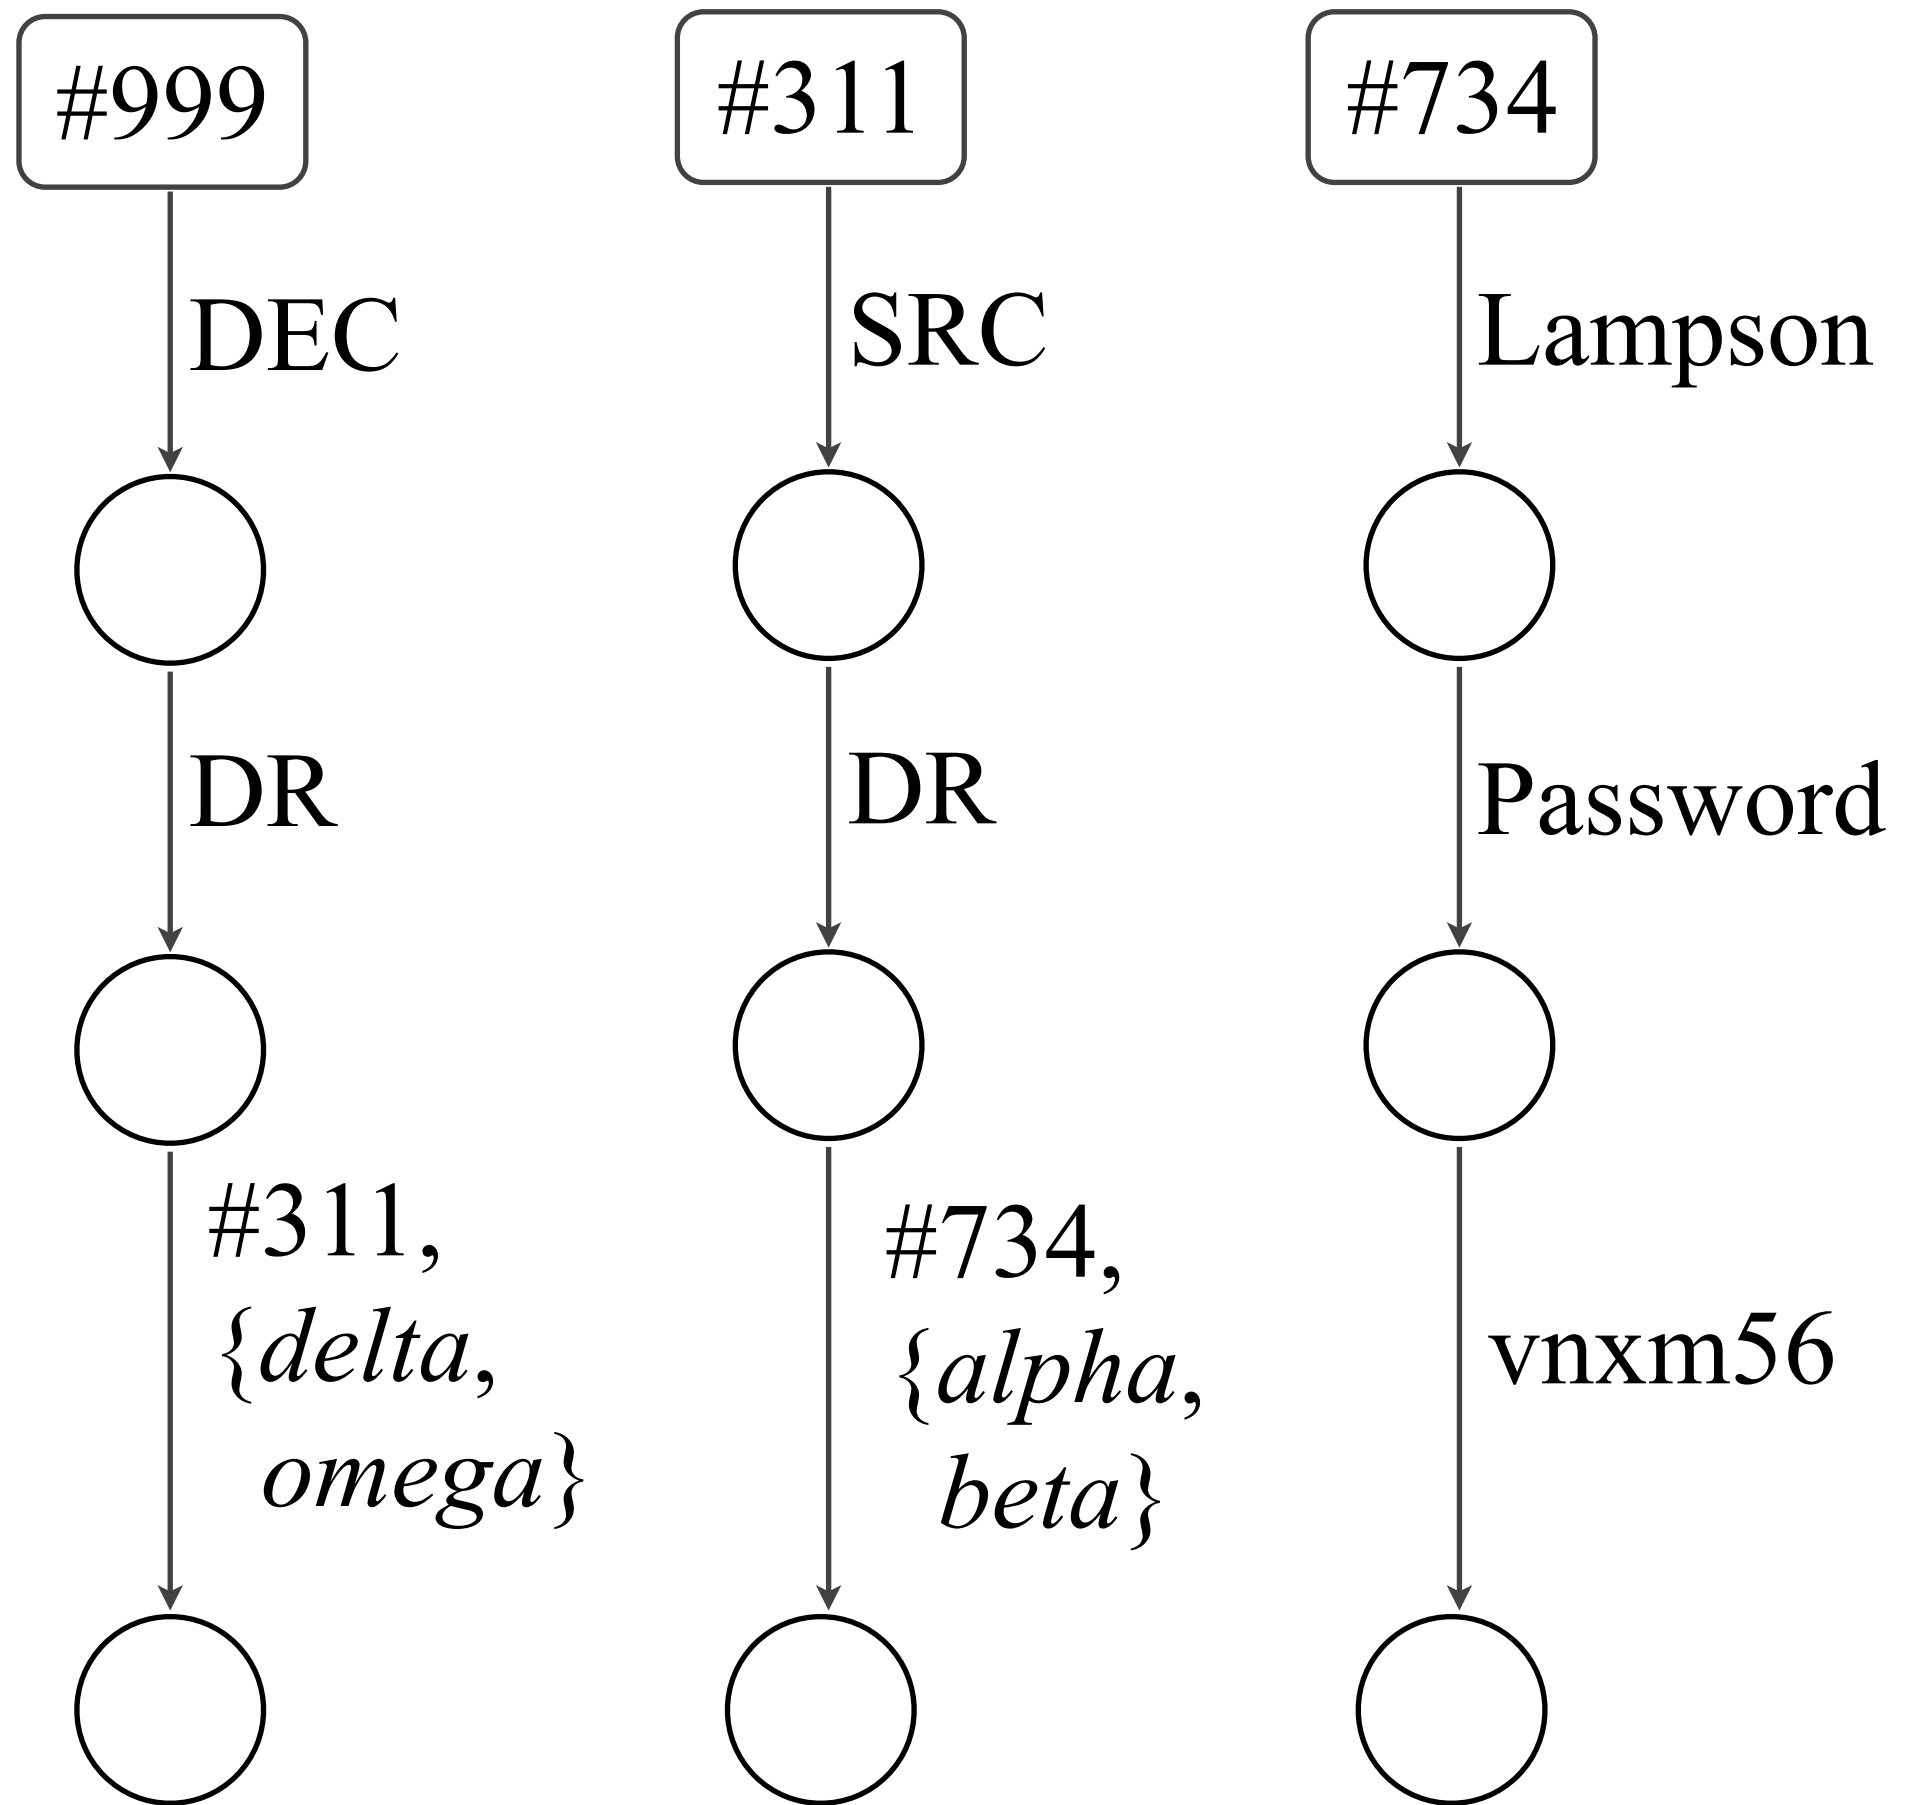
Names

Values

#999/DEC/SRC/Lampson/Password

#999/DEC/DR

#311, {*delta*, *omega*}



Names

Values

#999/DEC/SRC/Lampson/Password

#999/DEC/DR

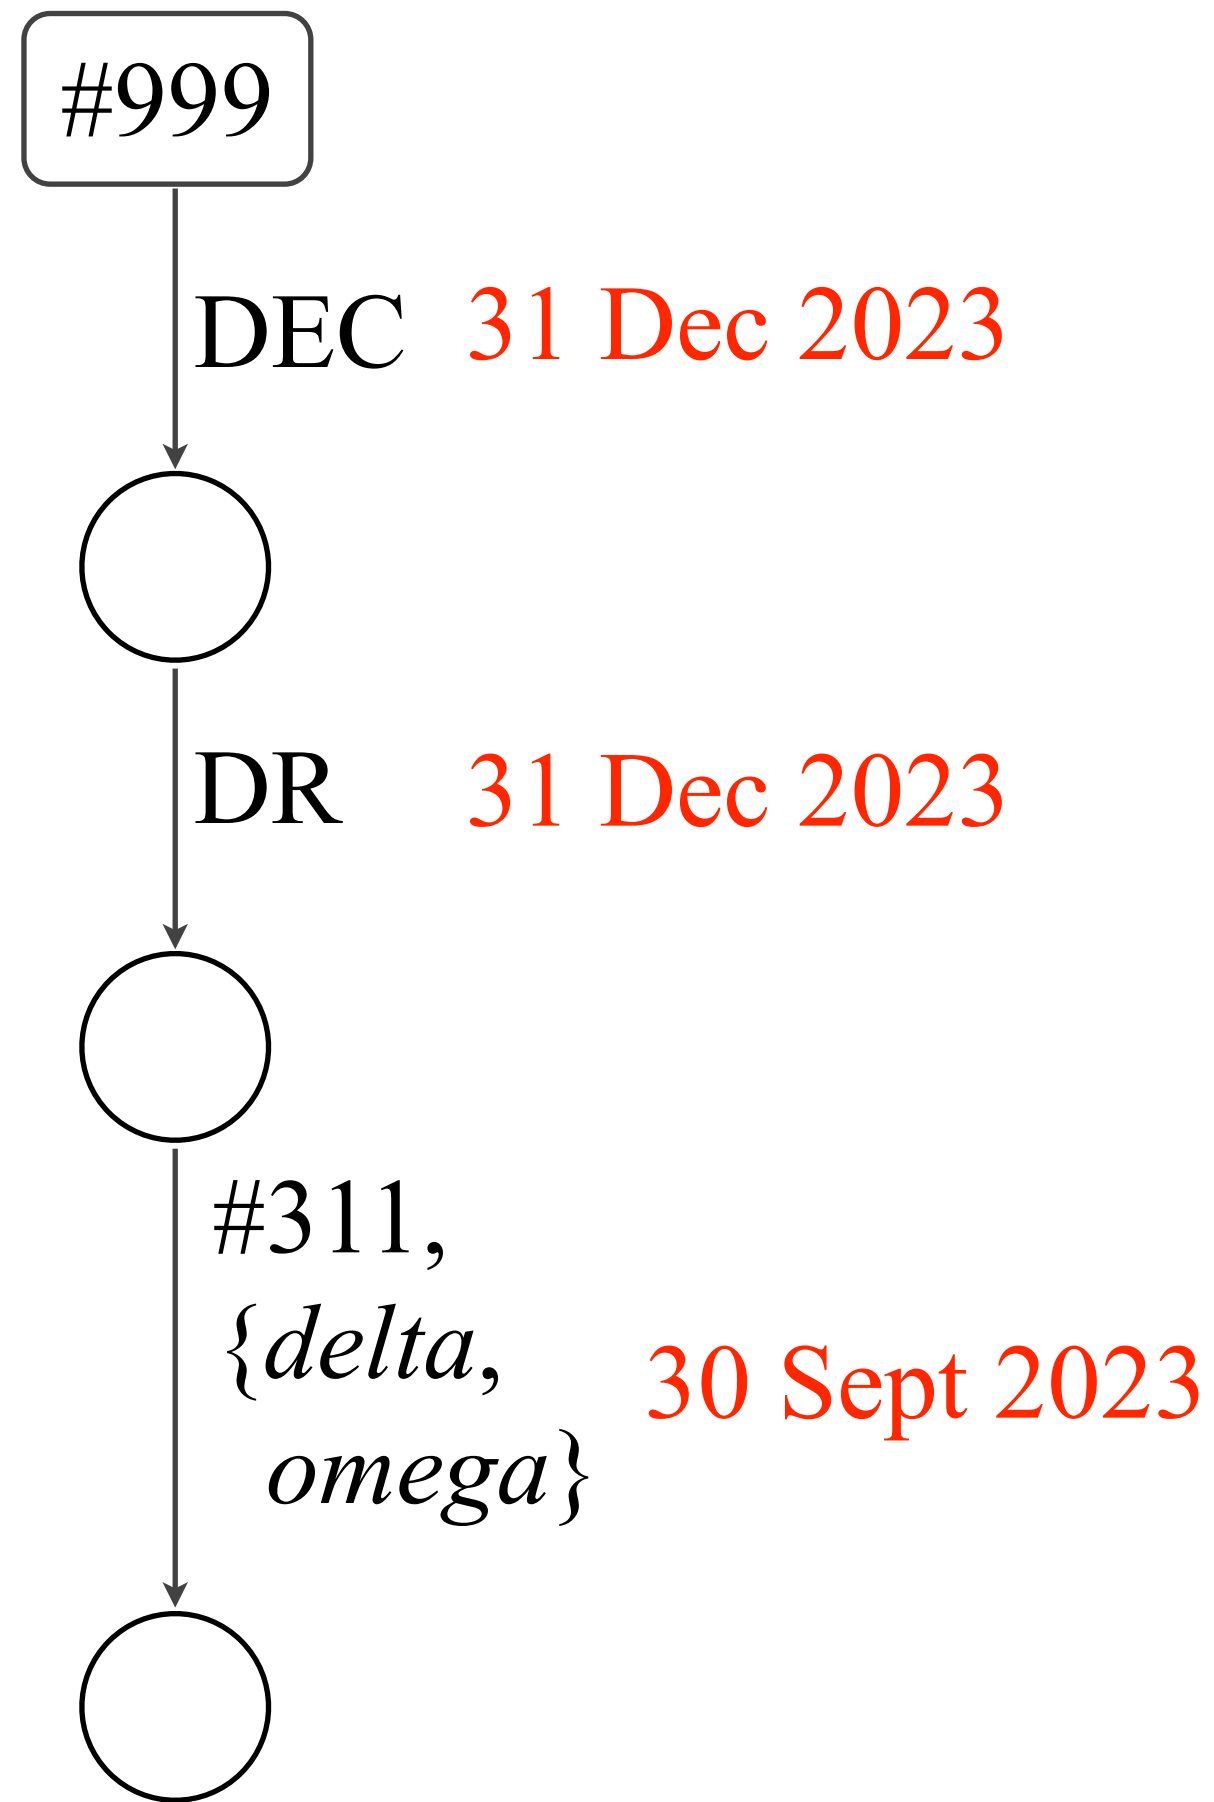
#311, {delta, omega}

#311/SRC/DR

#734, {alpha, beta}

#734/Lampson/Password

vnxm56



Names

Values

#999/DEC/SRC/Lampson/Password


#999/DEC/DR

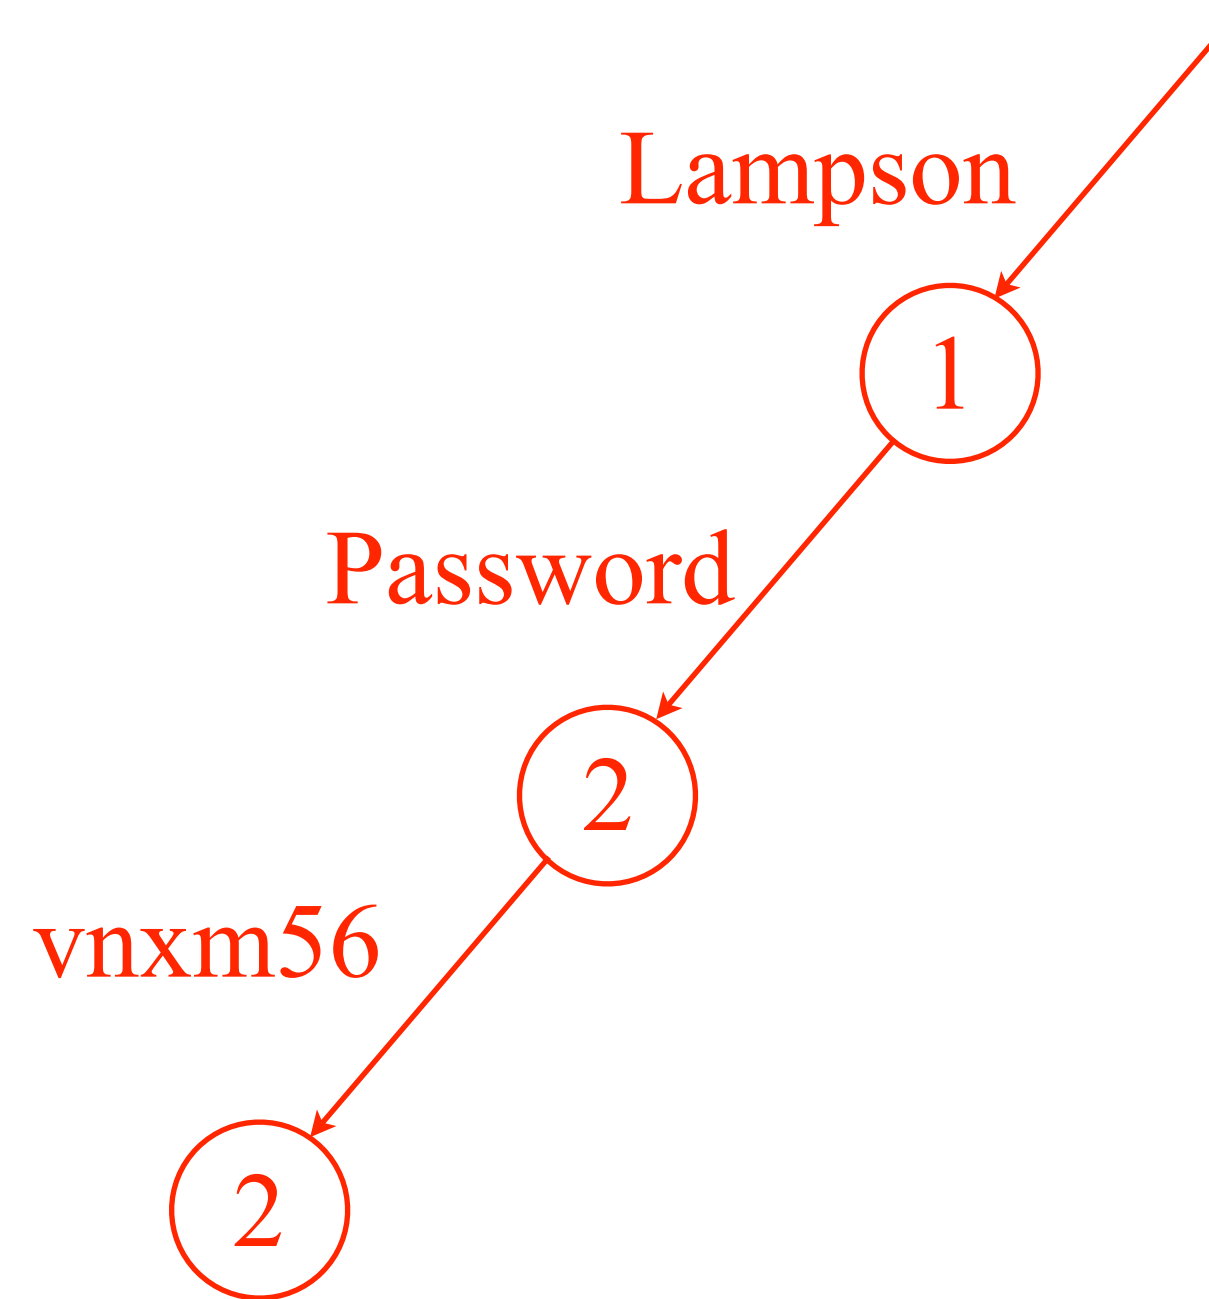

#311, {*delta*, *omega*}



Name lookup



- Clients can **cache** mappings to reduce latency
- Stale mappings avoided through **expiration times**

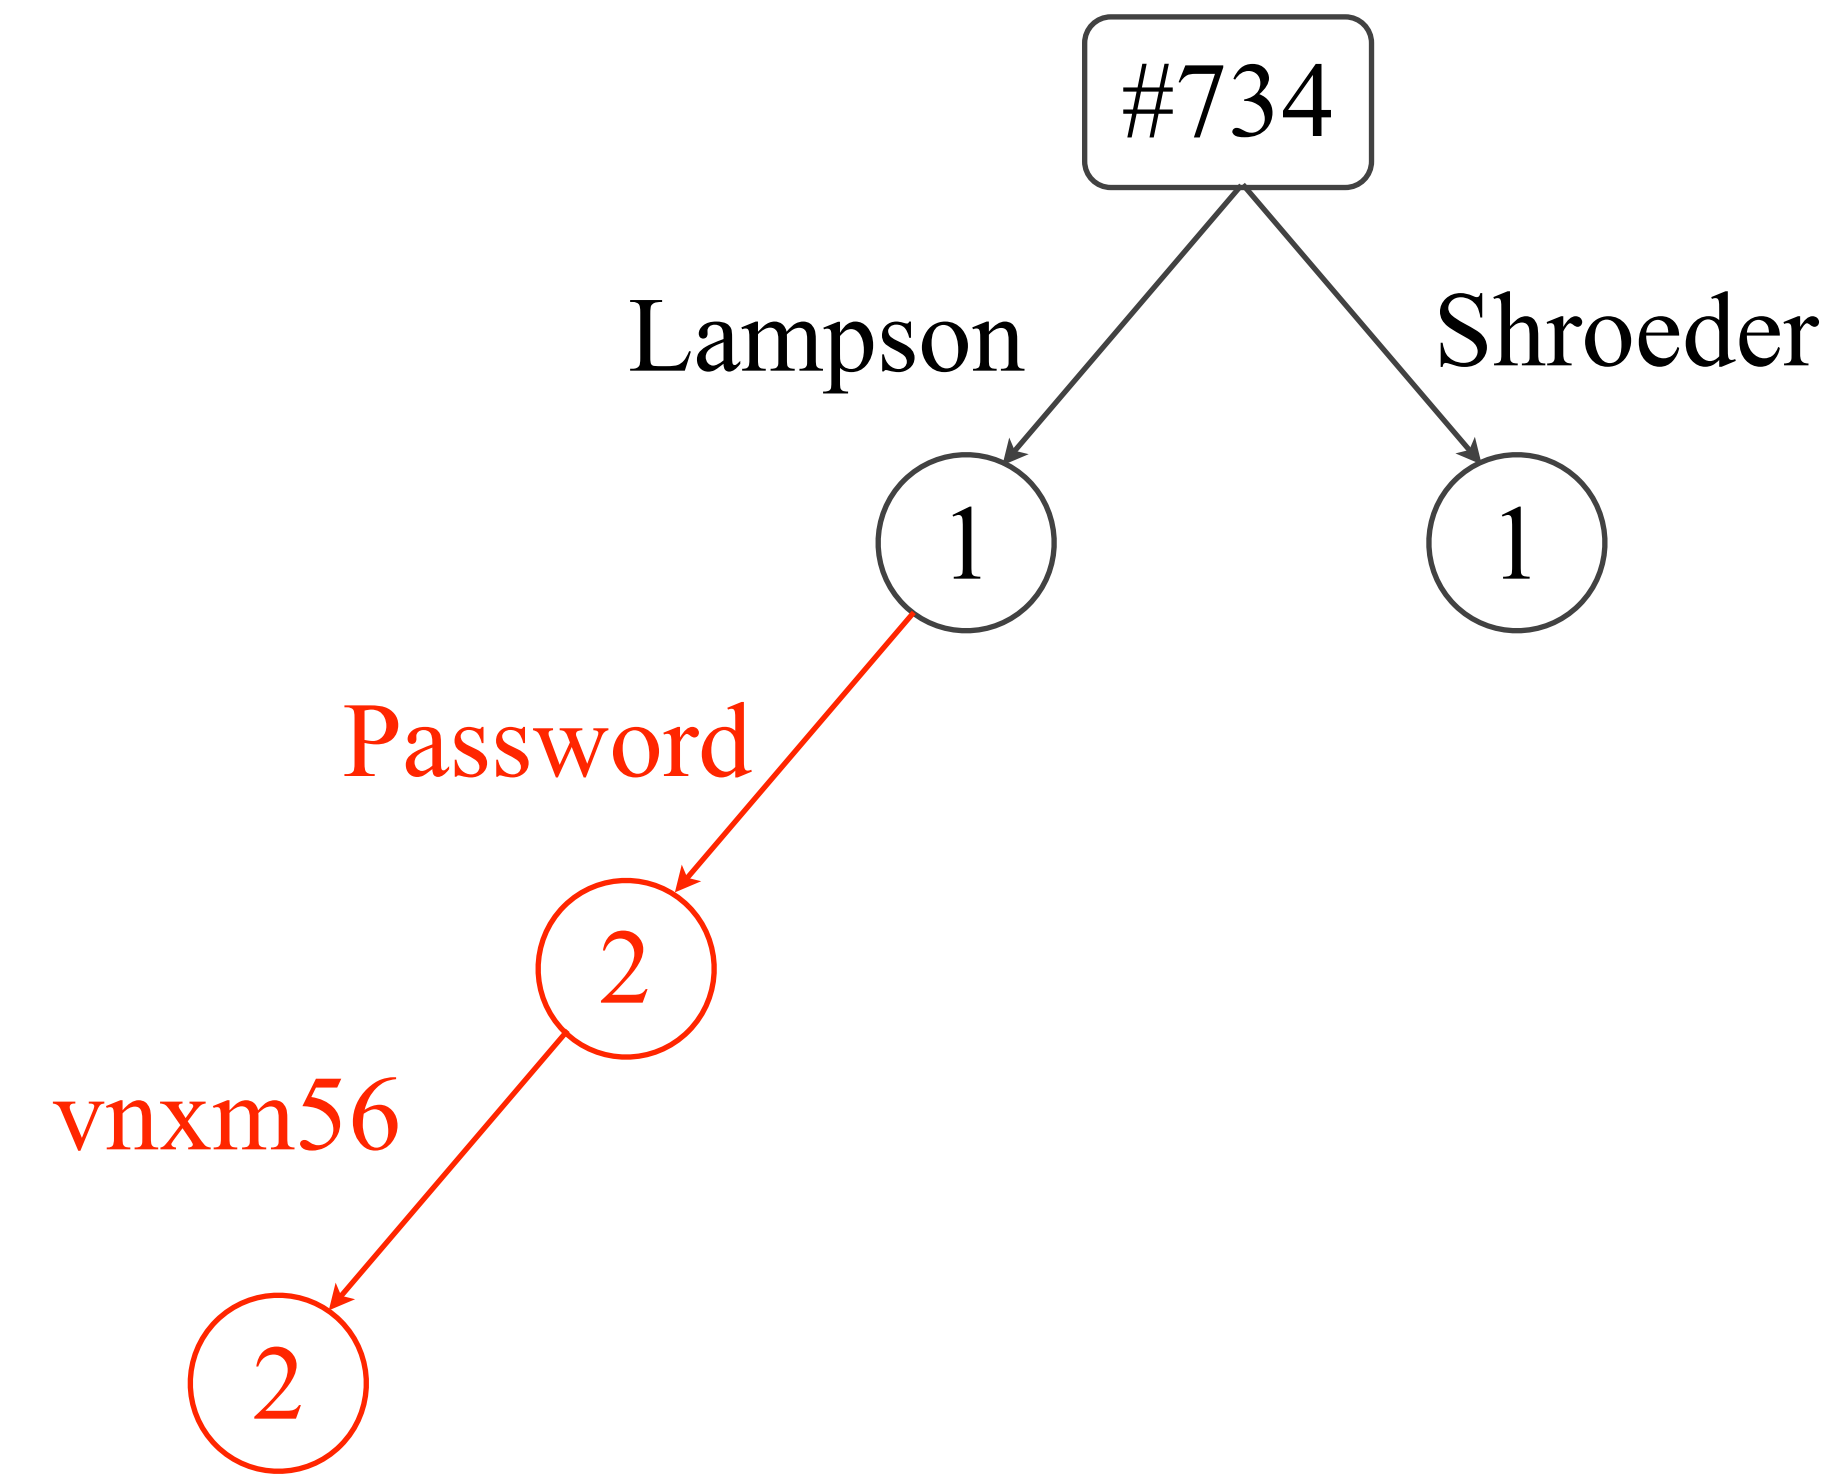
Updates






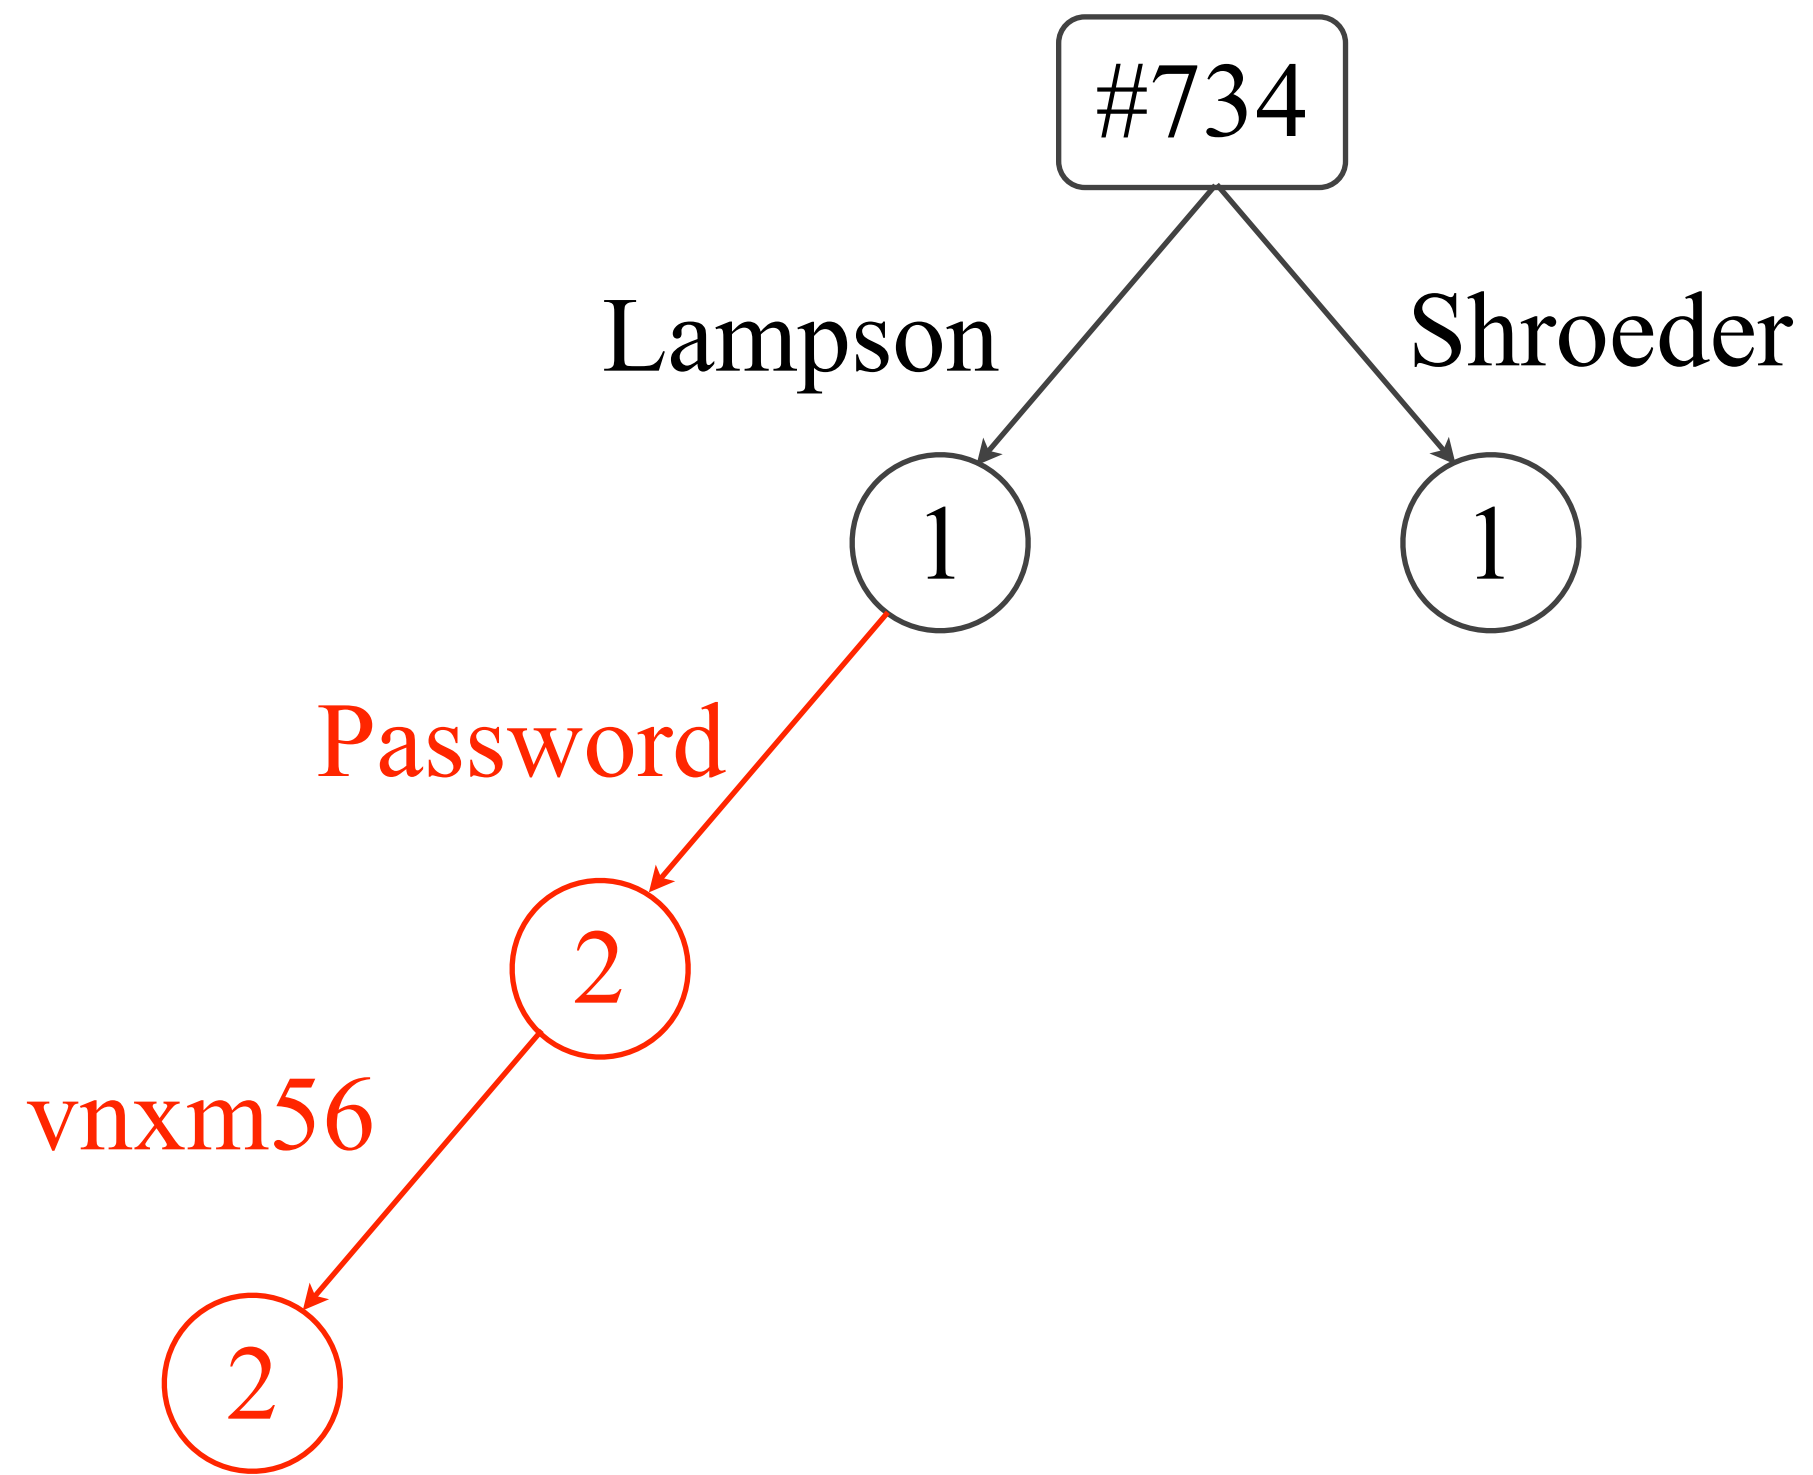














































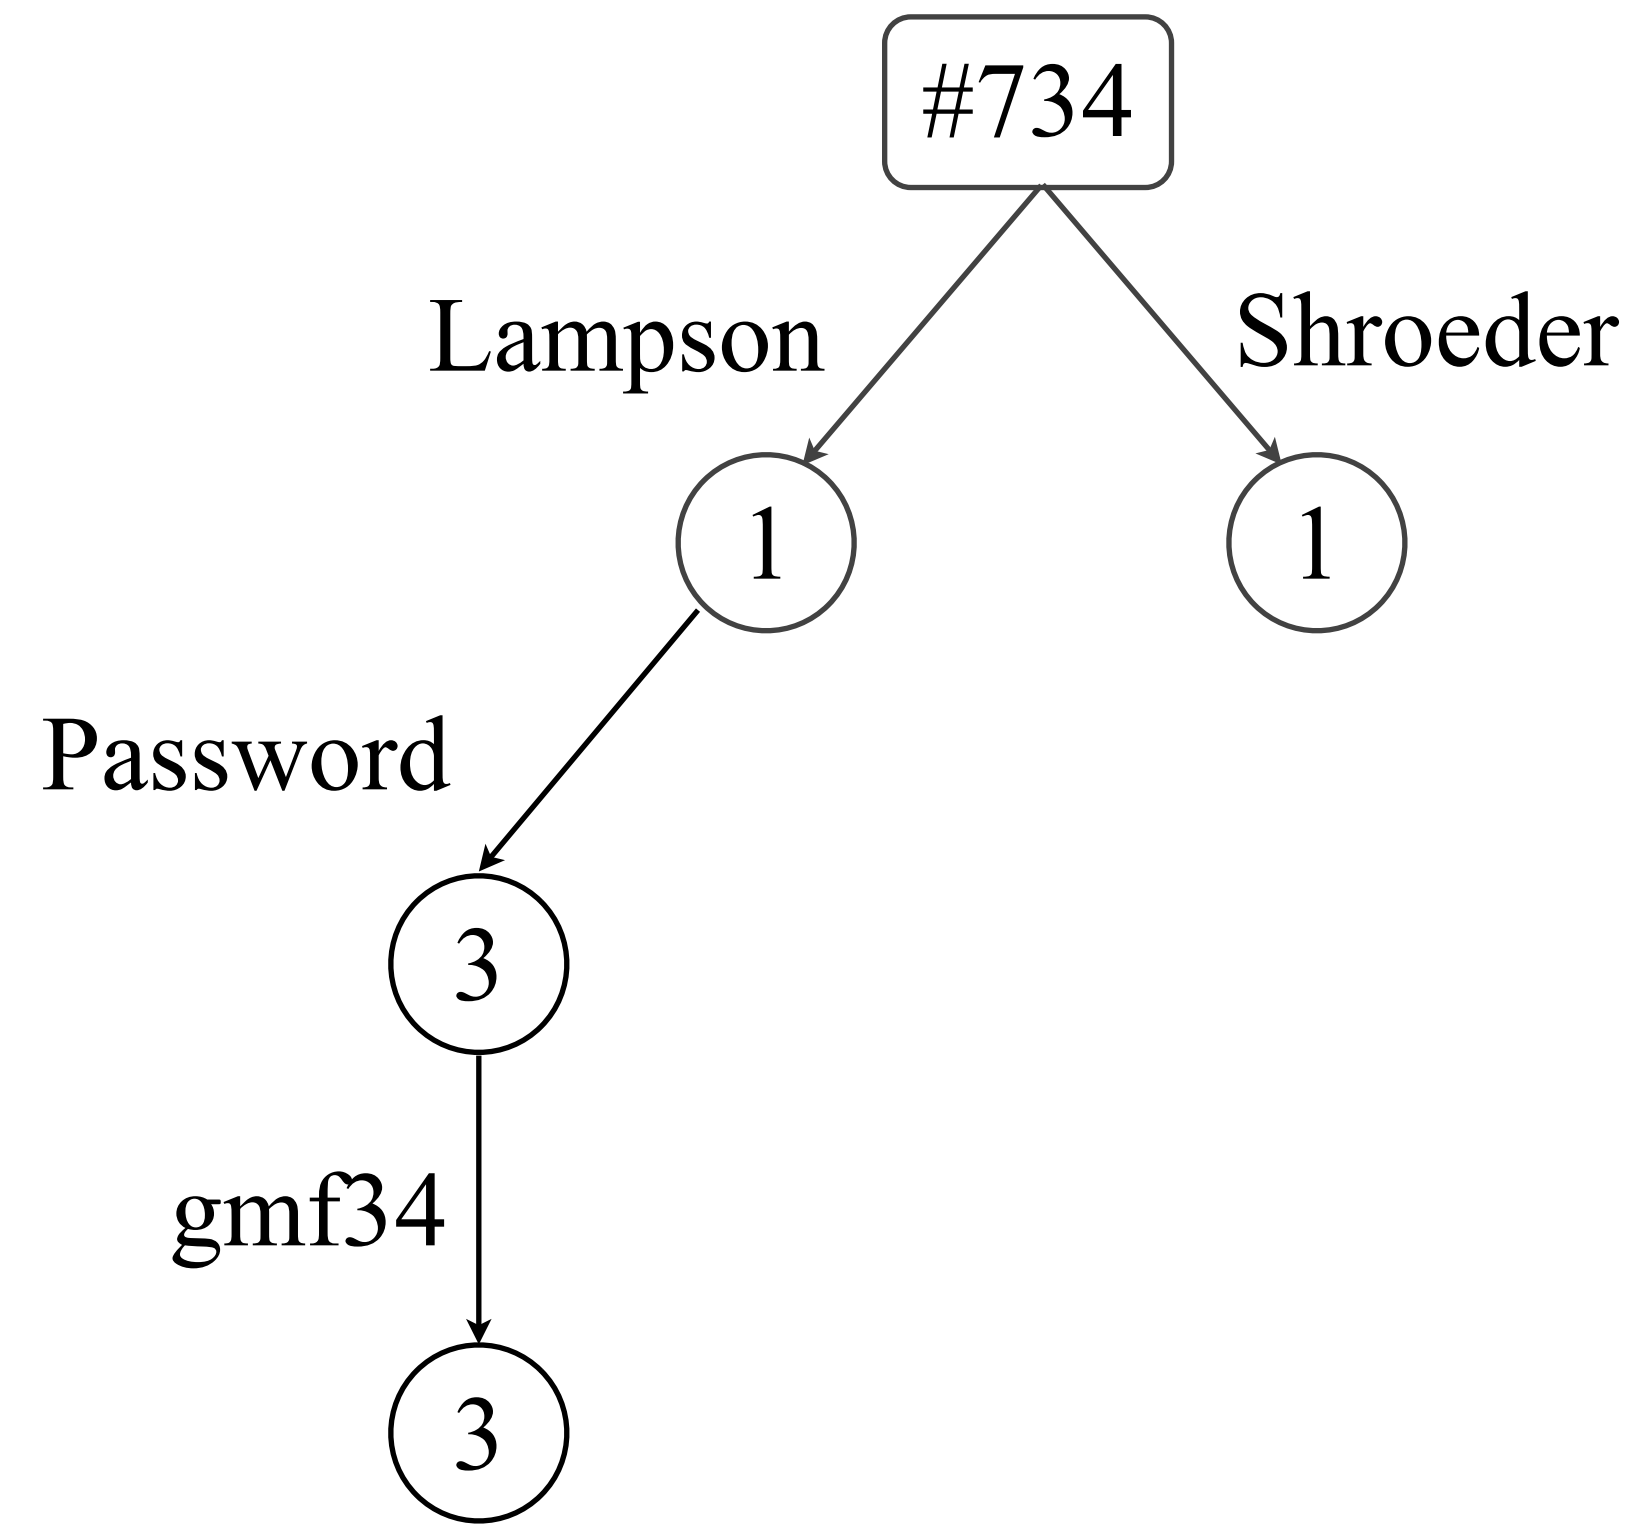


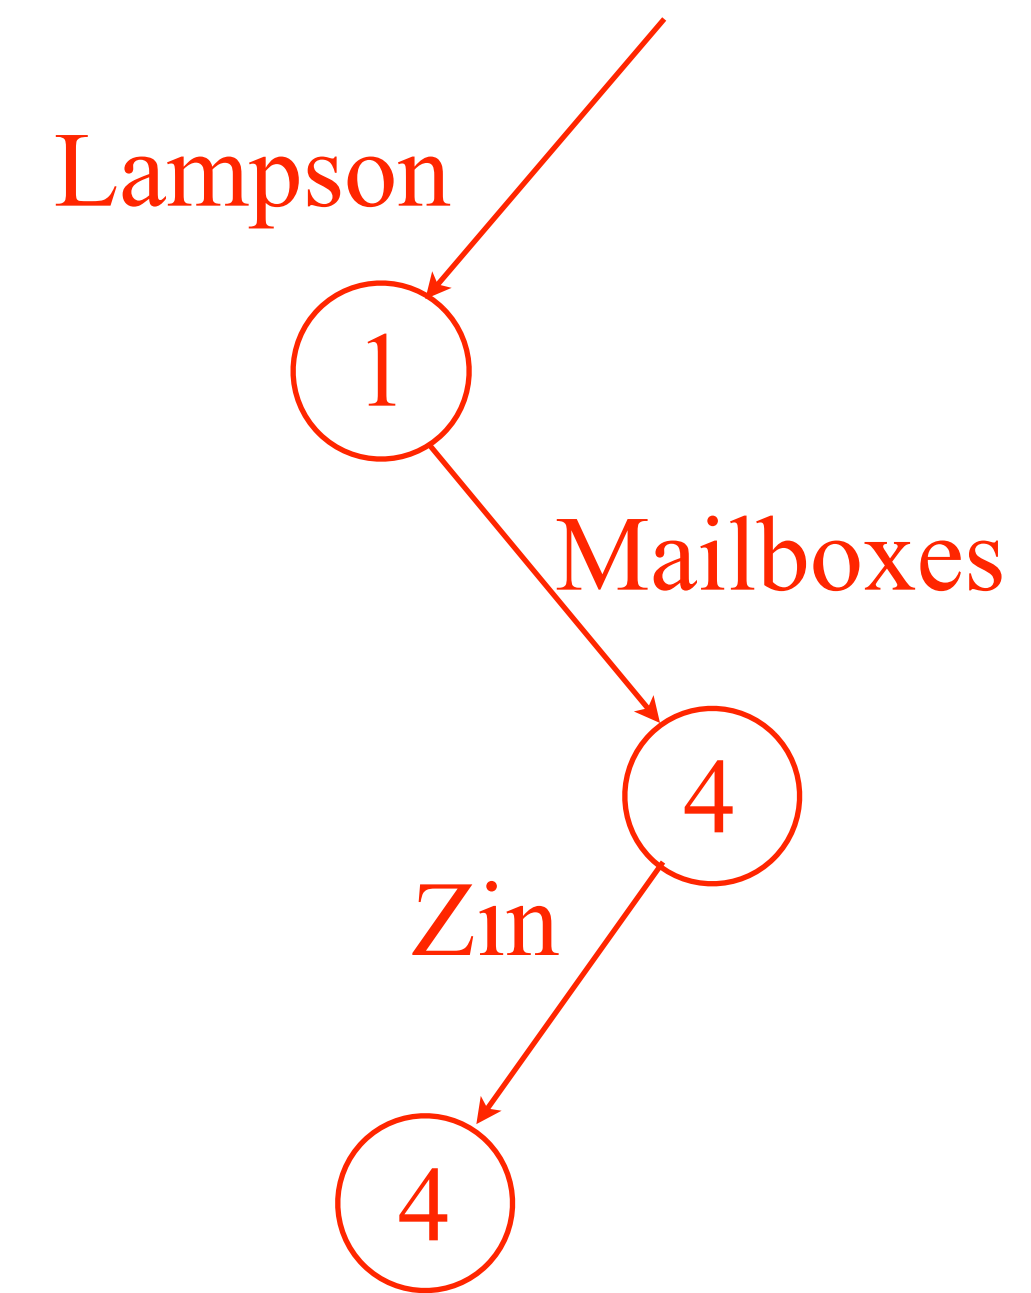



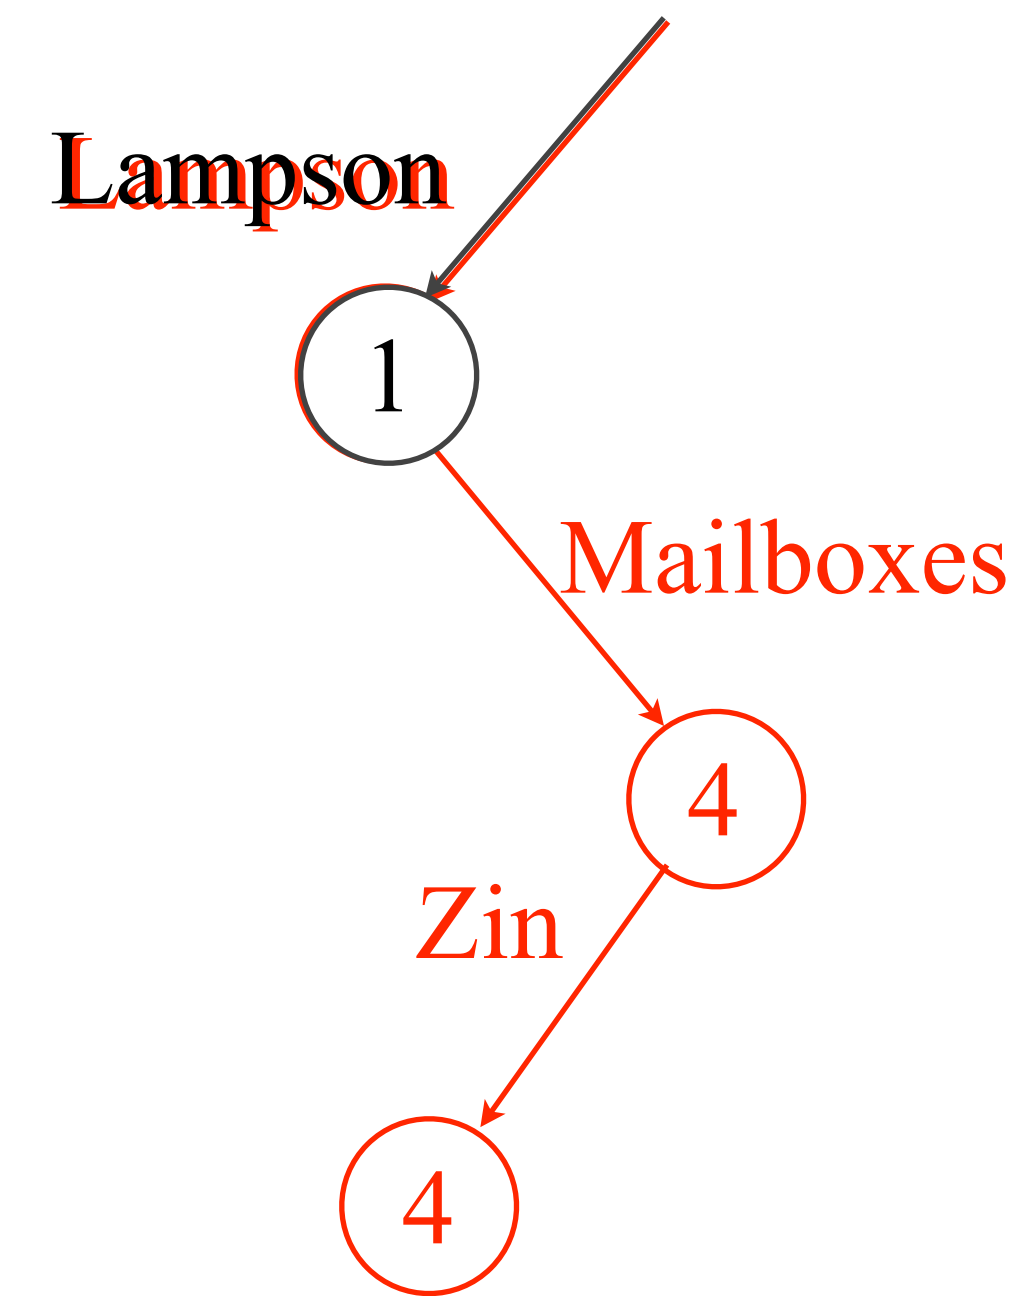
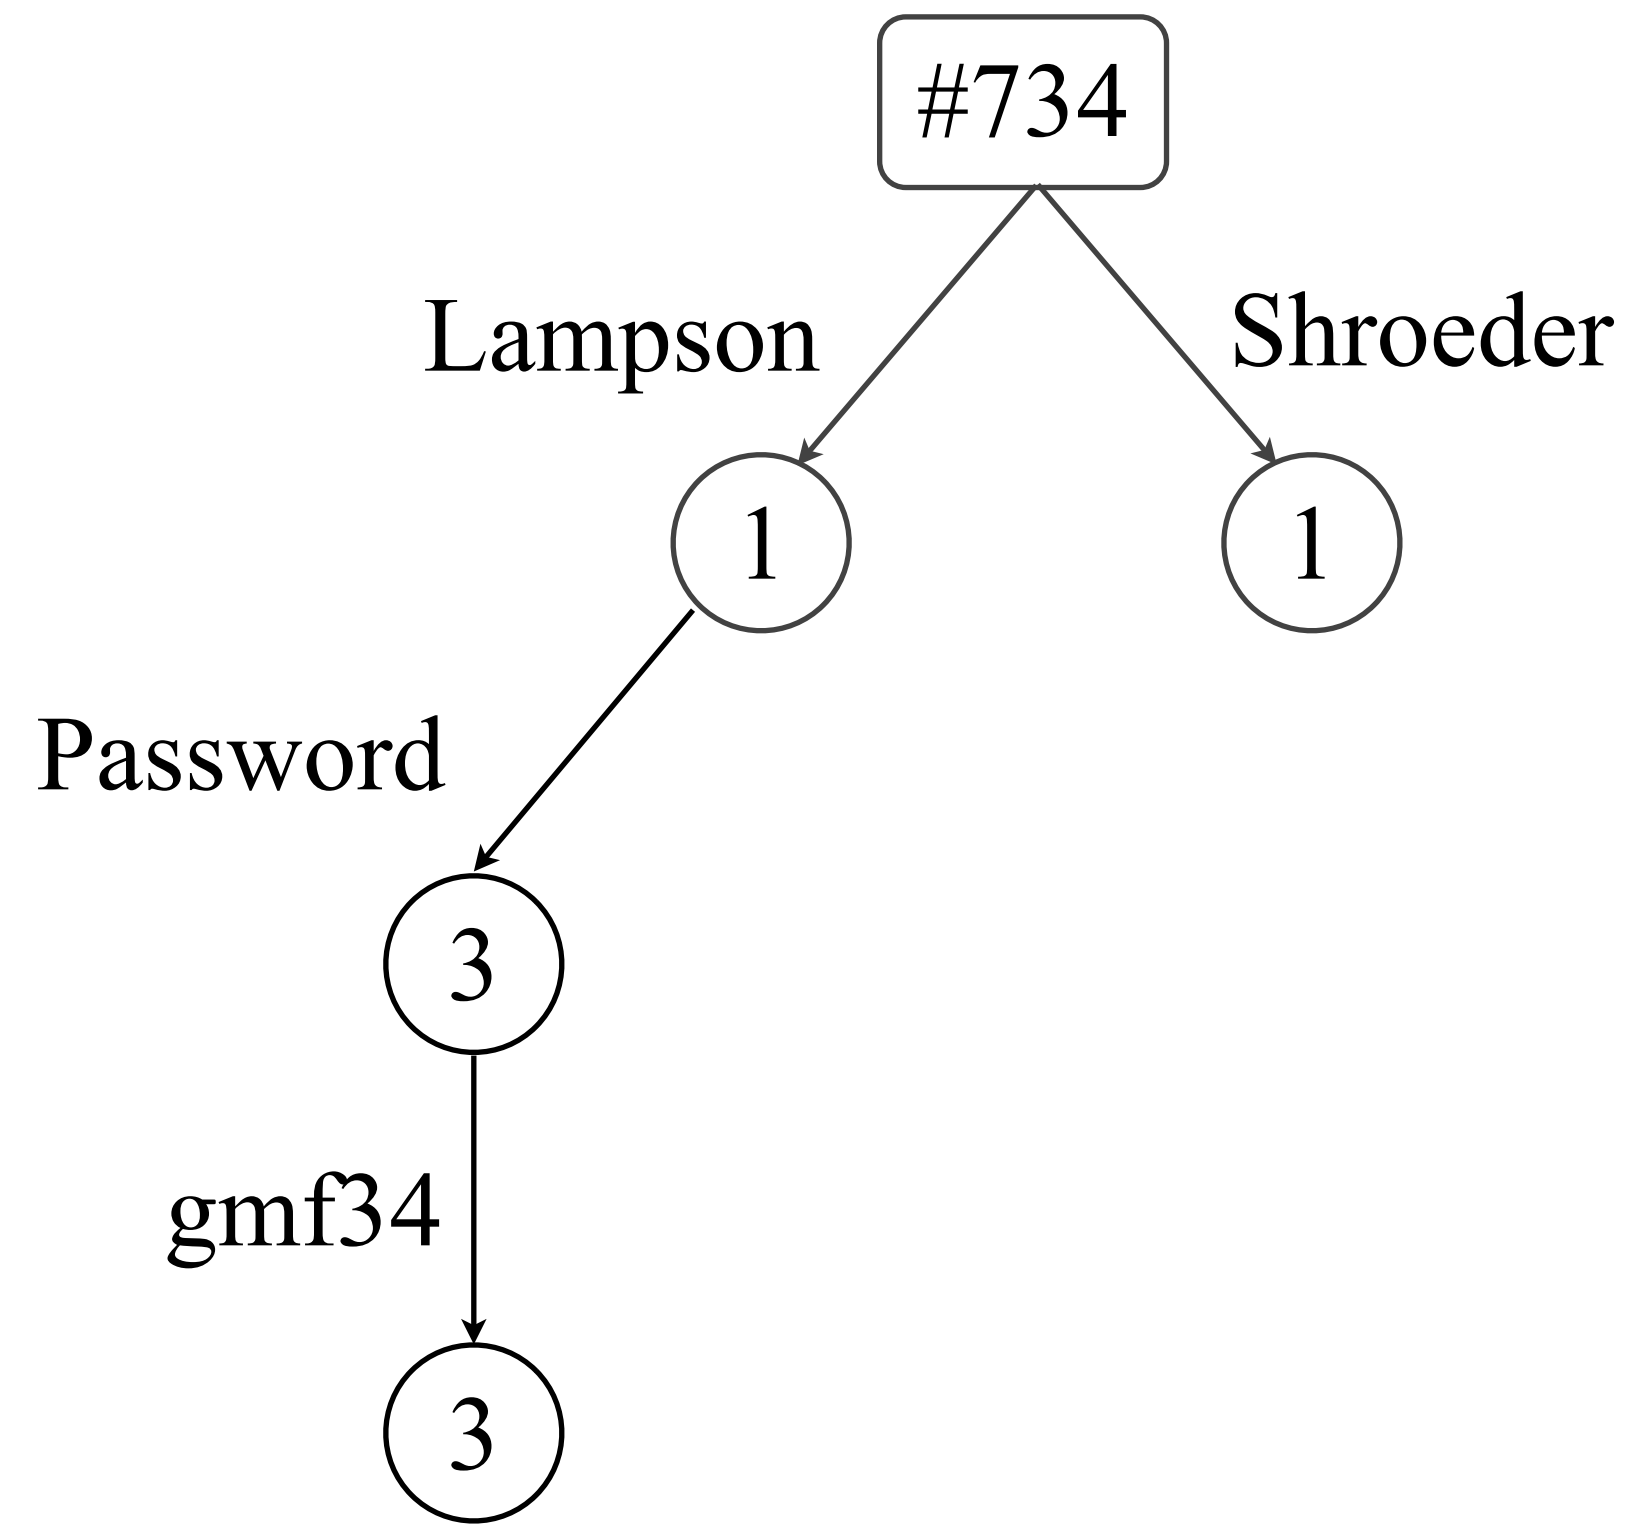


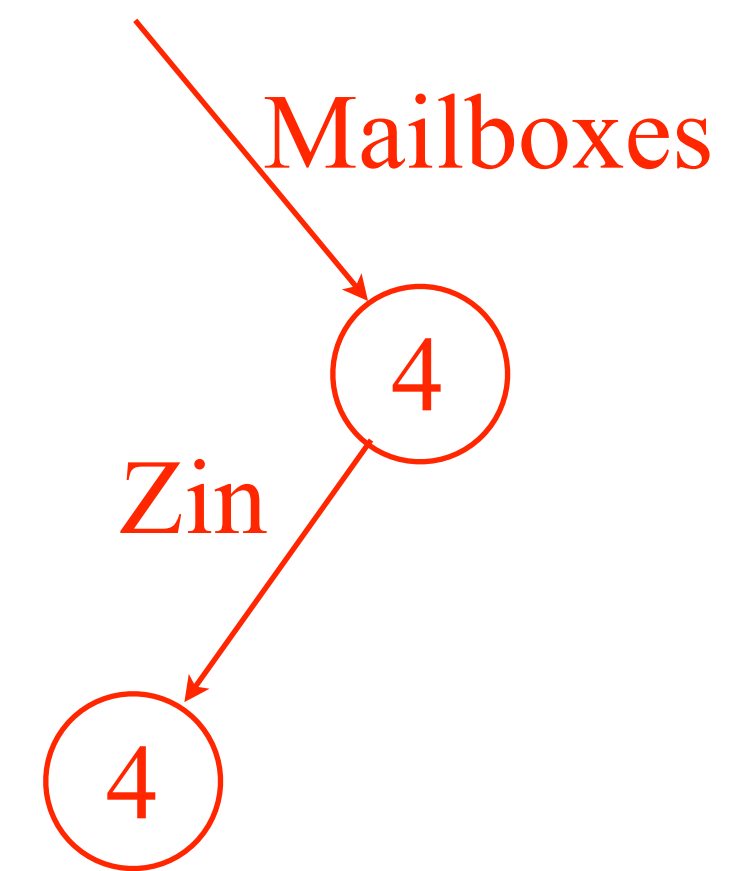



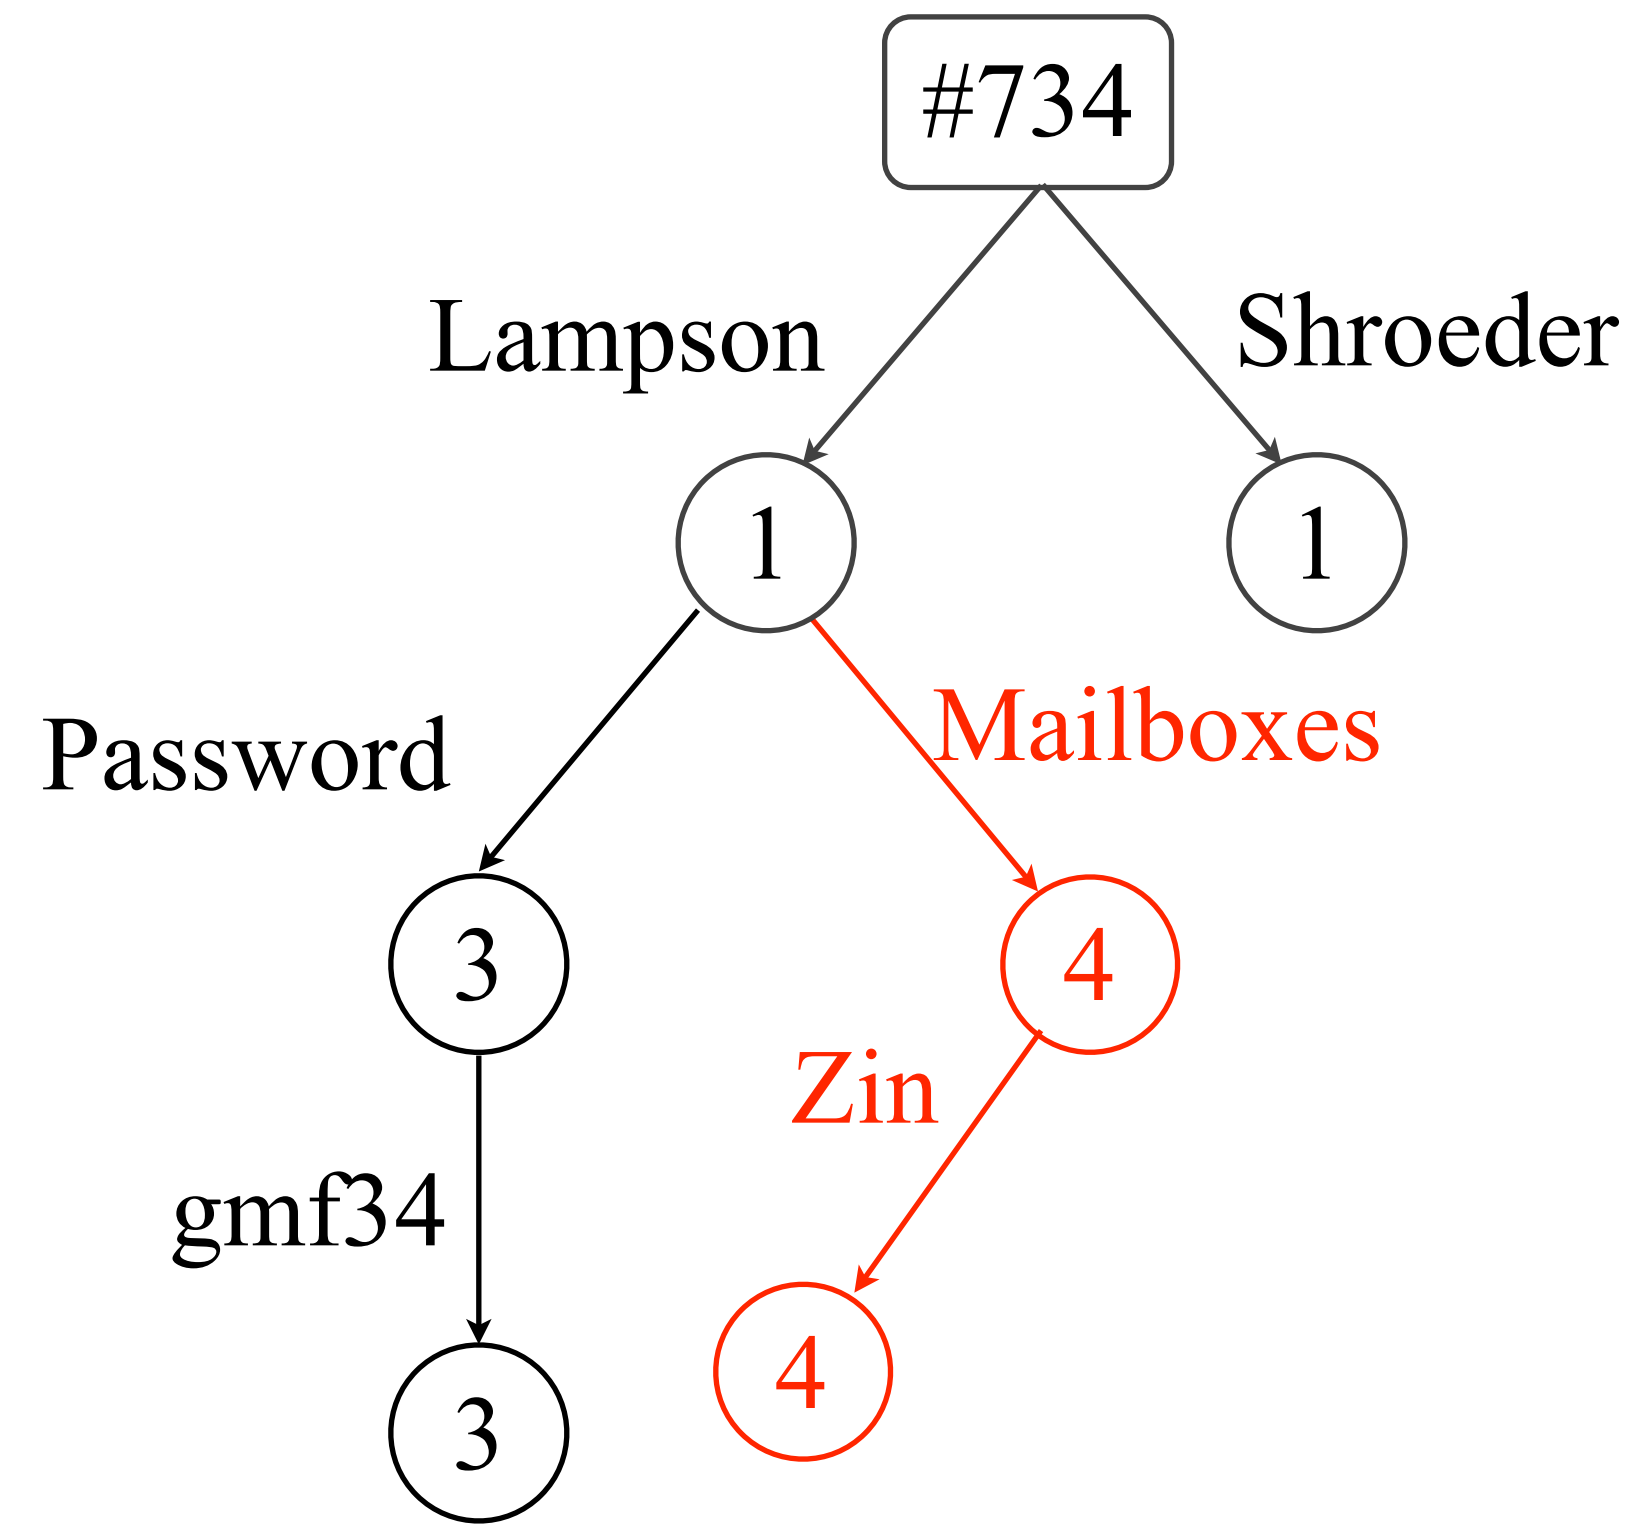










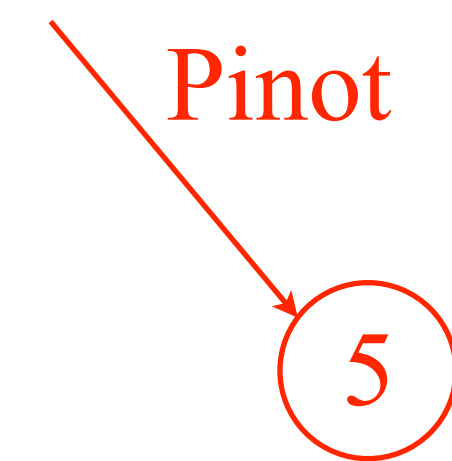
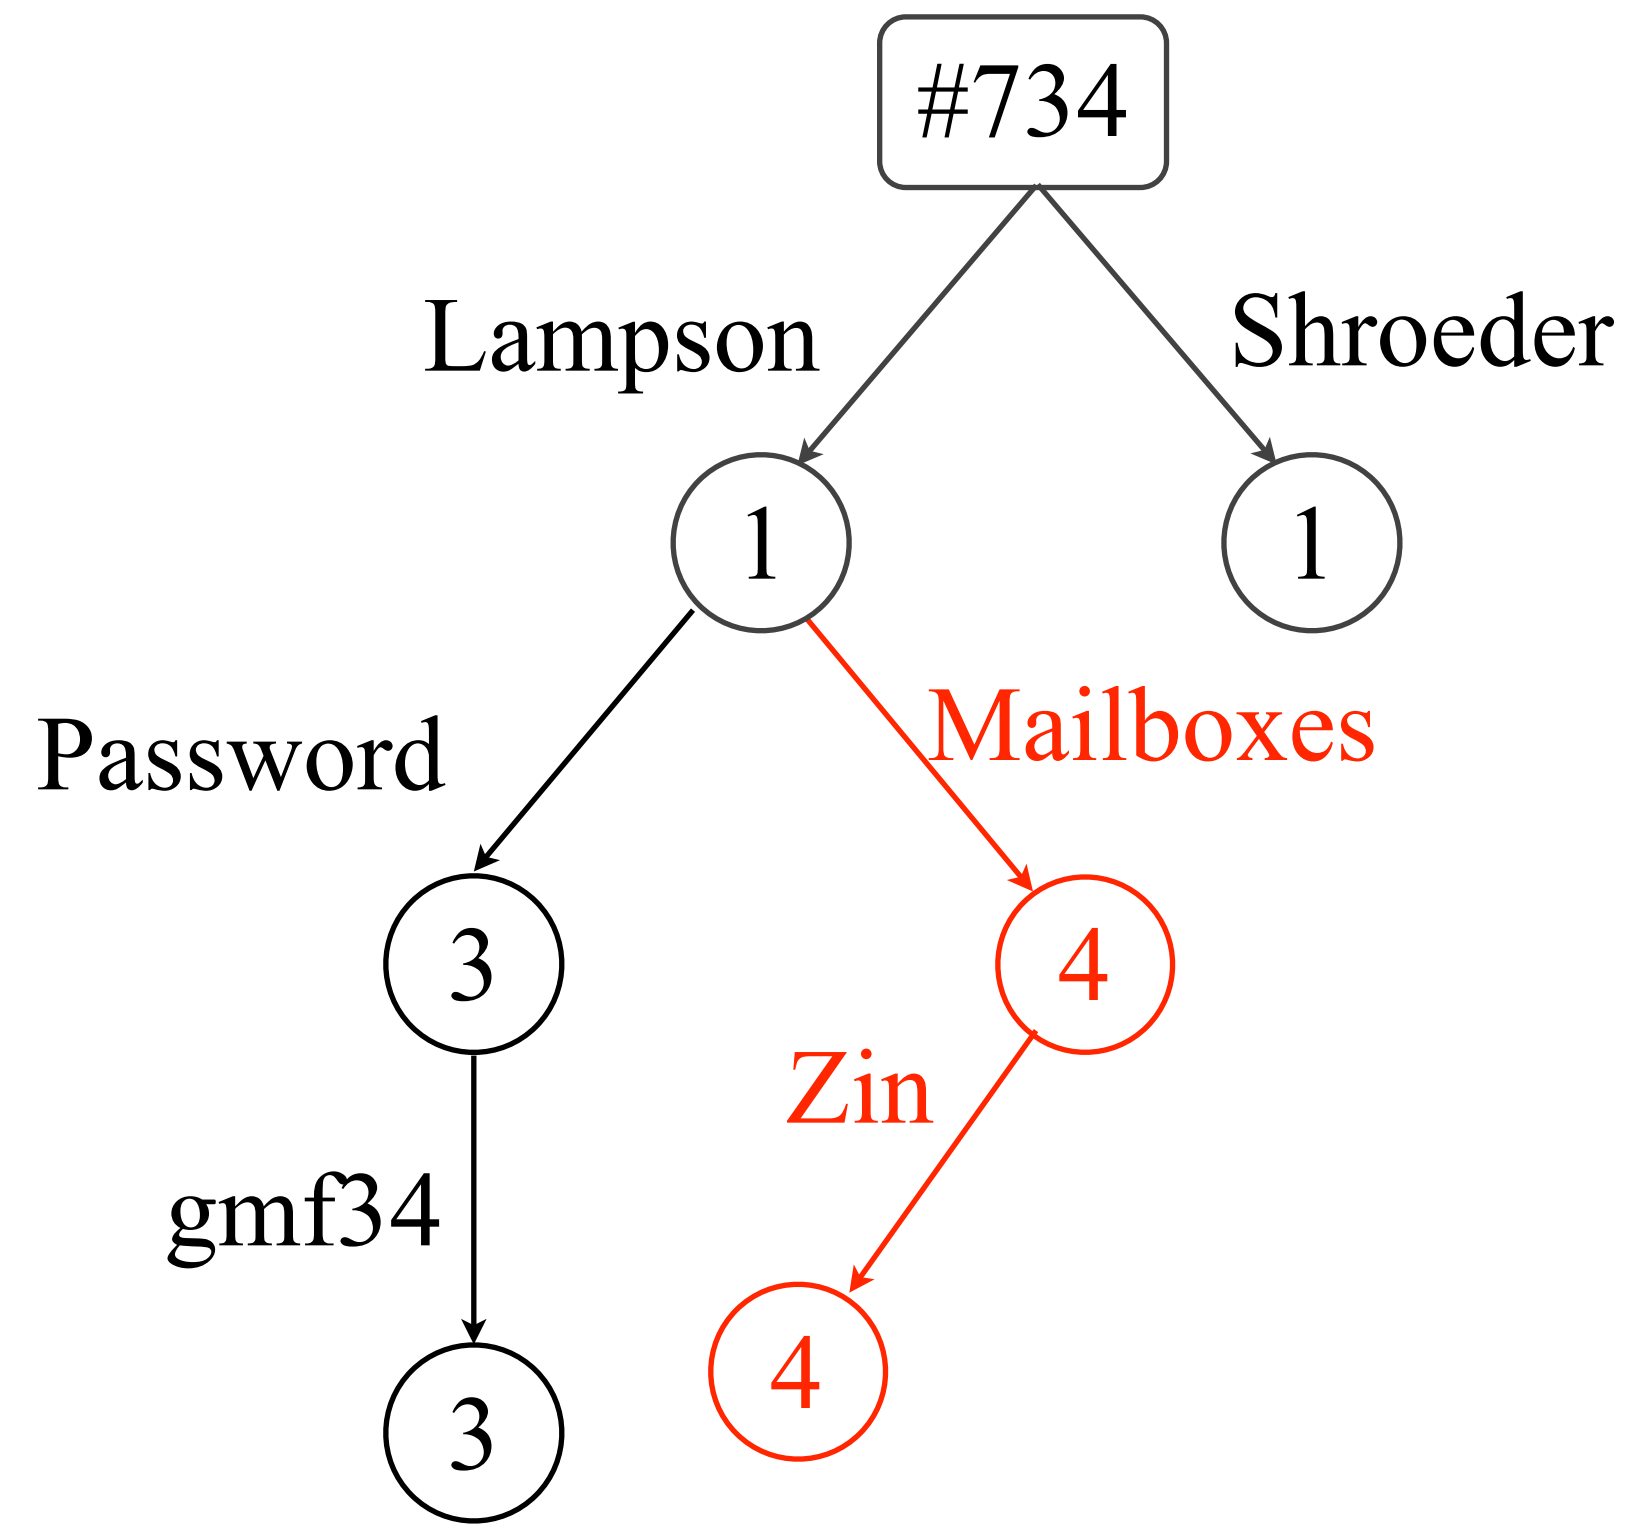


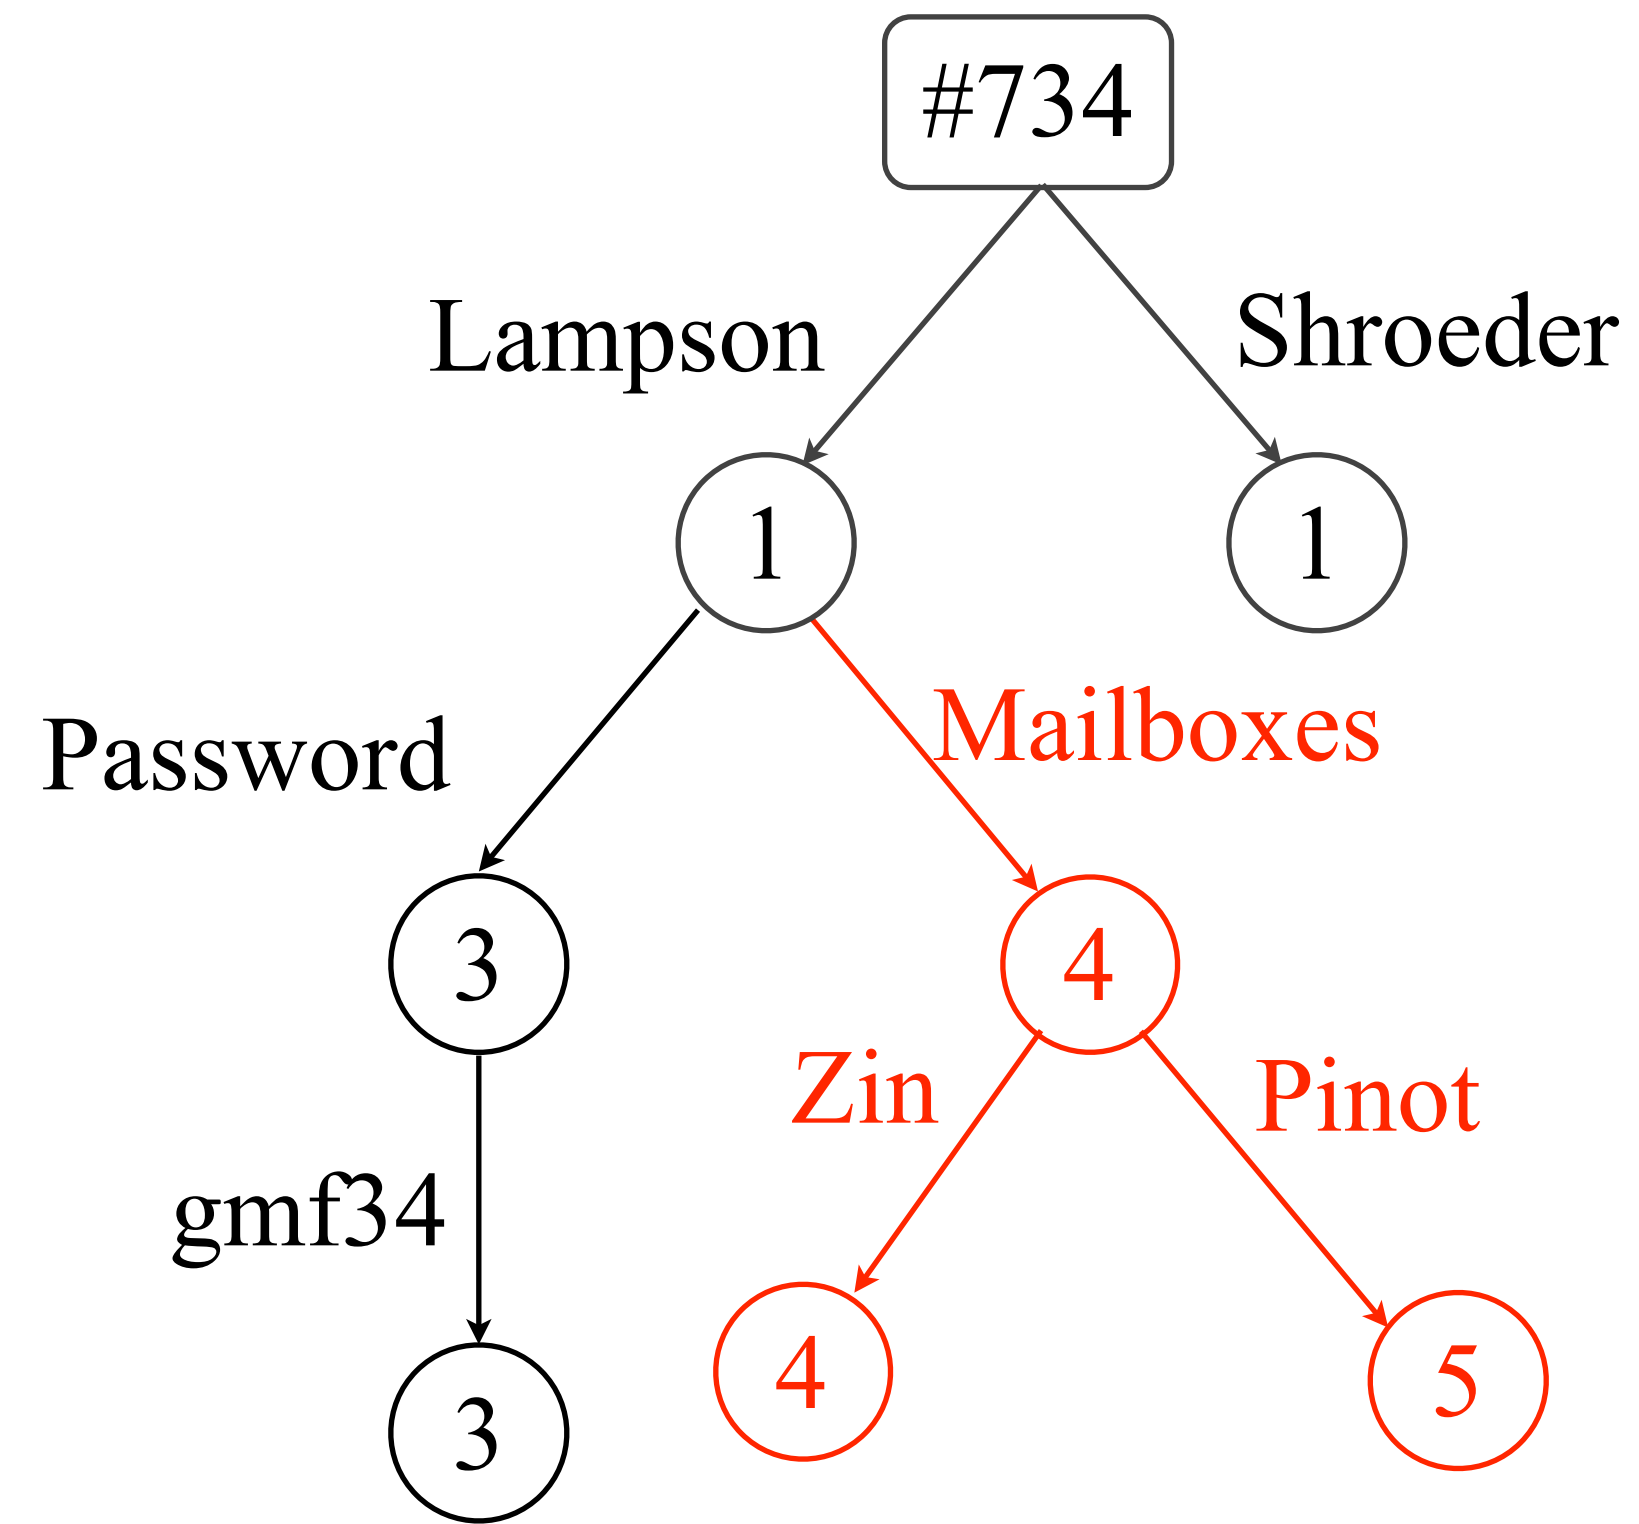







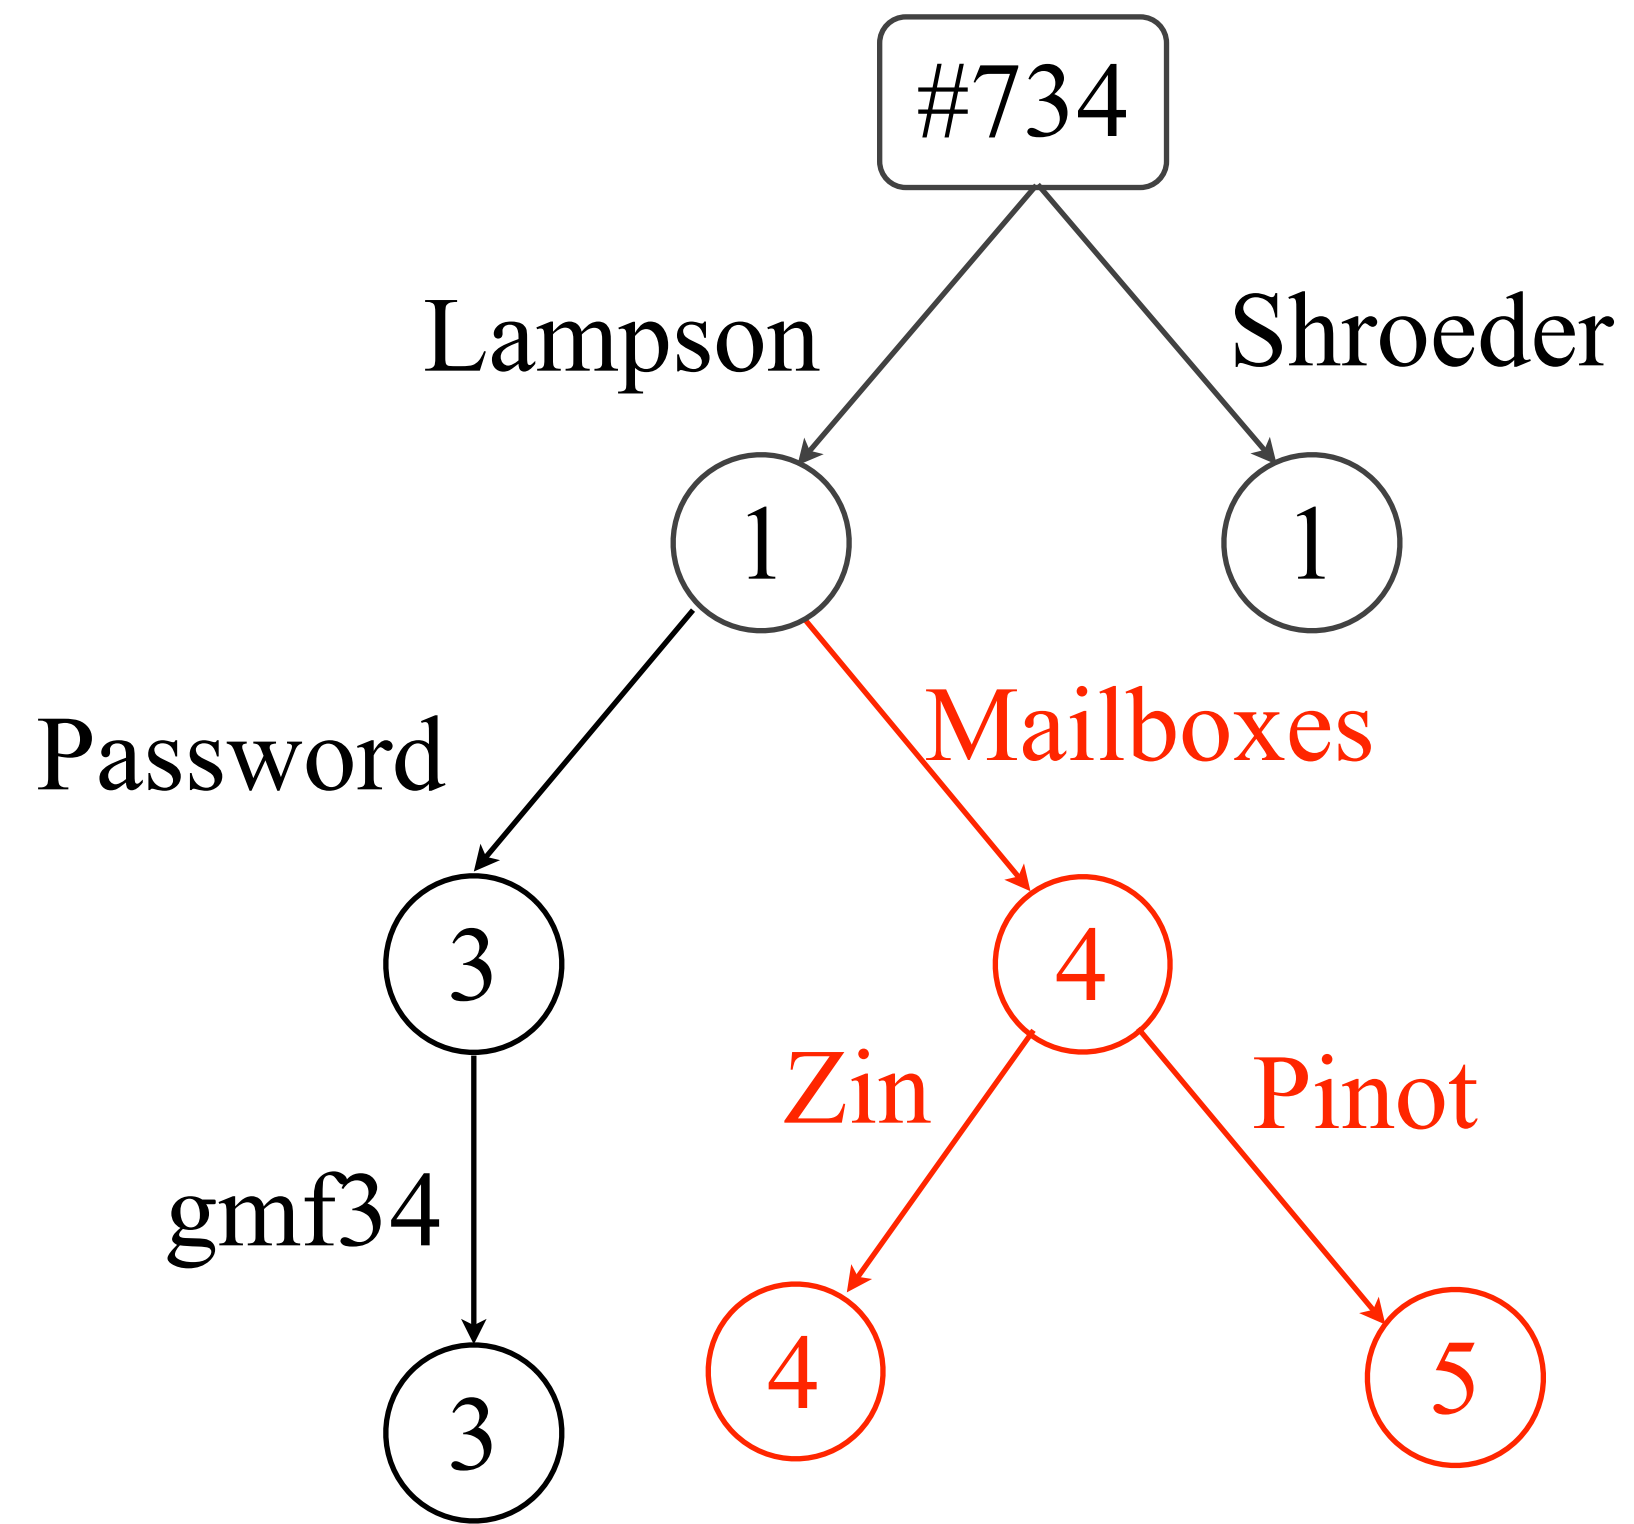


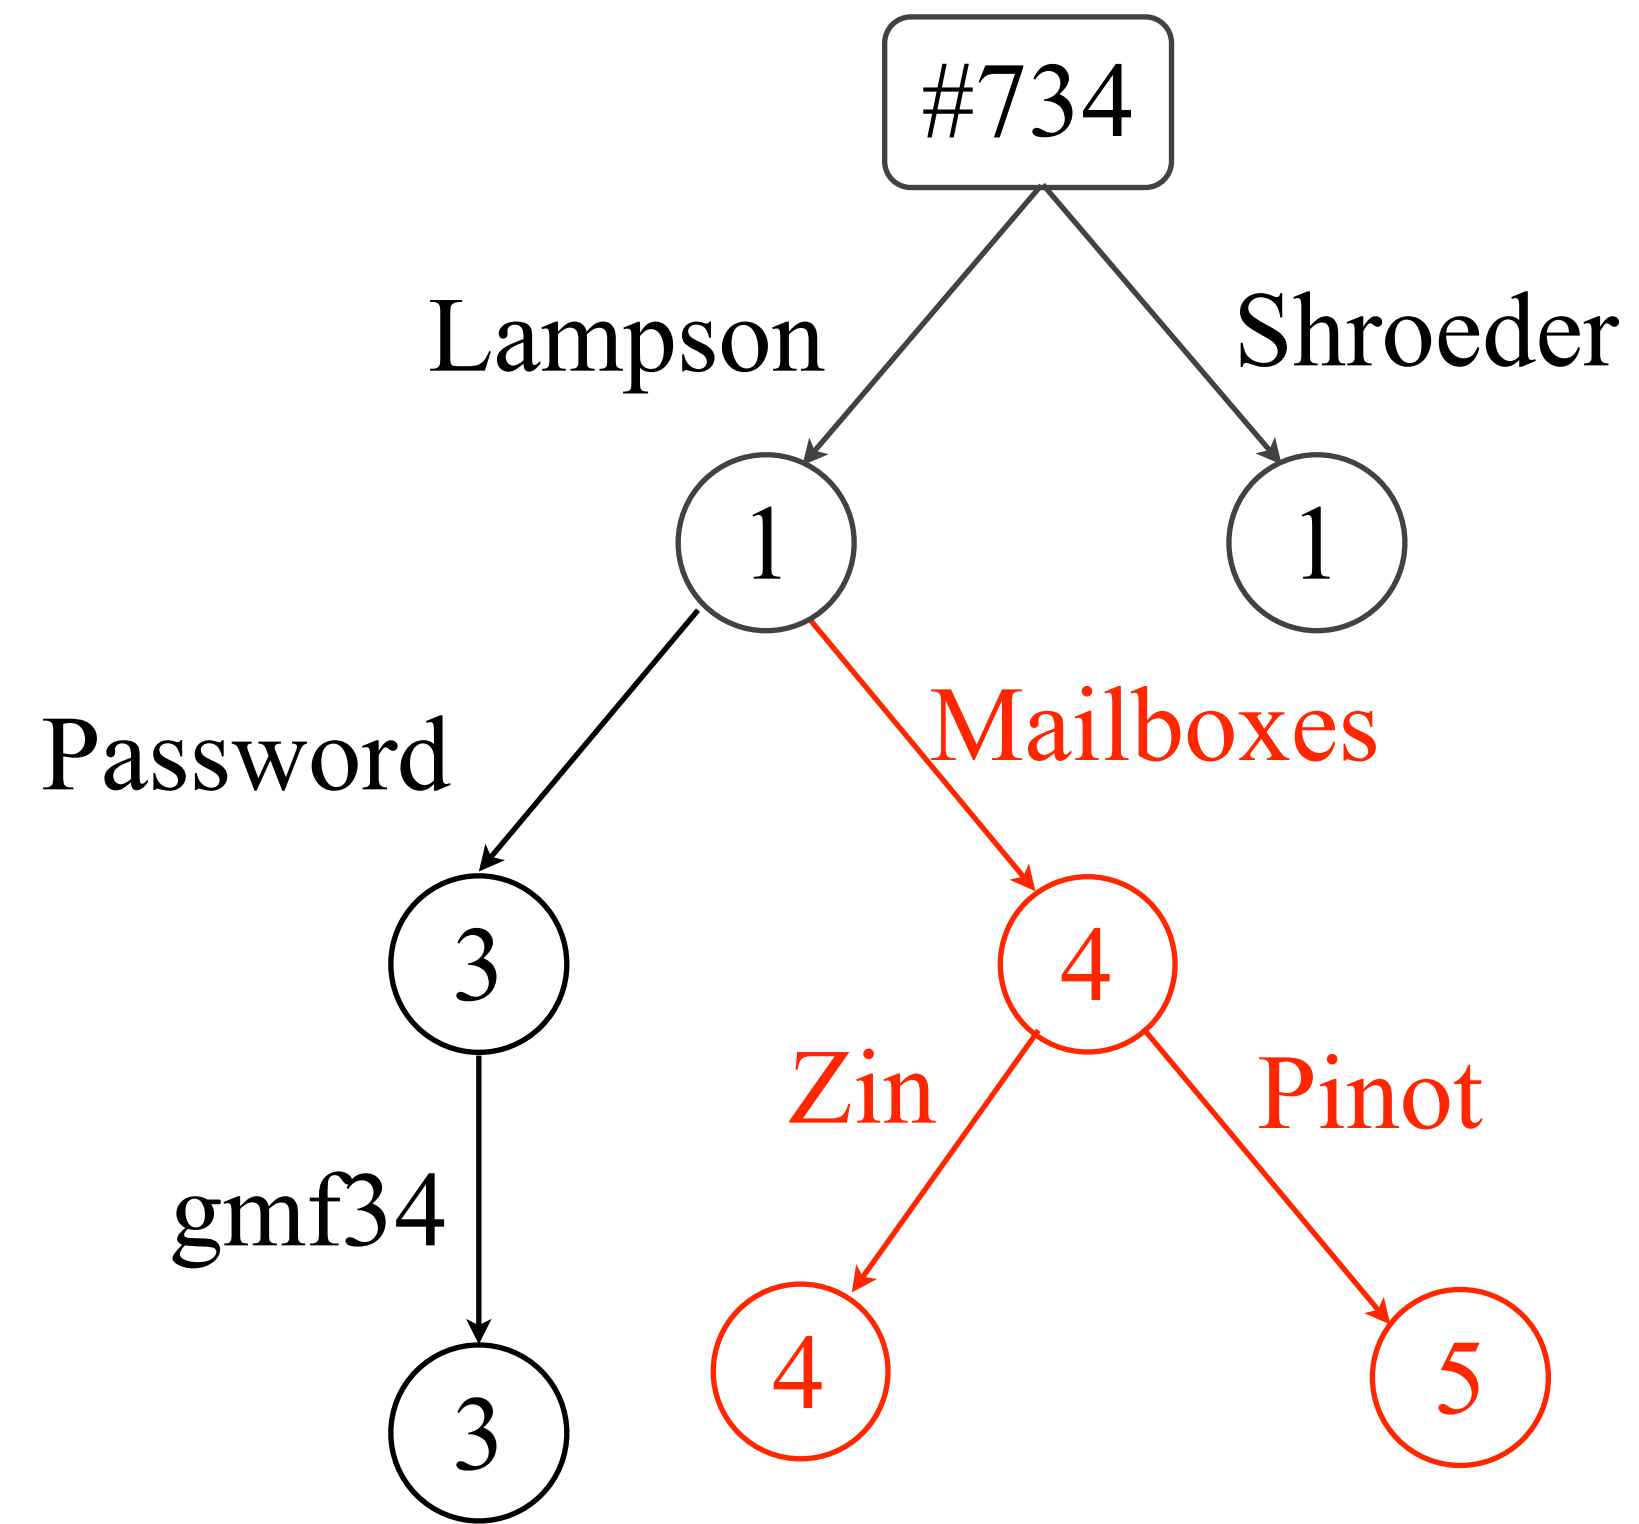












Updates

- **Commutative**: reordering updates does not affect the outcome
- **Idempotent**: reapplying updates does not affect the outcome
- Both achieved through timestamp ordering

Designing a Global Name Service

- **Scalability**: supports an arbitrary number of names and organizations
 - *achieved through hierarchy*
- **Fault-tolerance**: keeps functionality even when N servers fail
 - *achieved through redundancy + eventual consistency*
- Clients **cache** mappings
- Stale mappings avoided through **expiration dates**
- Updates are **commutative** and **idempotent**