5617, Spring 2020
computer networking and communication

anduo wang, Temple University

TTLMAN 302, R 17:30-20:00

A Data-Oriented (and Beyond) Network Architecture

https://dl.acm.org/doi/10.1145/1282380.1282402

DNS — a historical perspective

fundamental part of today's Internet

- underlying almost all network usage

ad-hoc design

- created late, basic pieces of the Internet already in place
  - TCP bound to IP addresses
- limit the extent to which DNS could permeate the Internet architecture

DNS — a historical perspective

fundamental part of today's Internet

- underlying almost all network usage

ad-hoc design

- created late, basic pieces of the Internet already in place
  - TCP bound to IP addresses
- limit the extent to which DNS could permeate the Internet architecture

this paper

a clean slate, principled redesign of DNS services

motivation

(broad) goal: a network architecture gracefully accommodates a wide spectrum of uses

- in the present and for the future

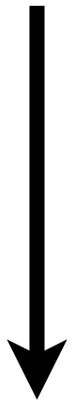

motivation

(broad) goal: a network architecture gracefully accommodates a wide spectrum of uses

- in the present and for the future

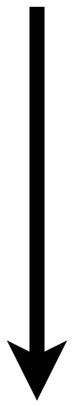

(specific) goal: address the shift in usage from host-centric to data-centric applications

data centric application

user cares about content, not location (of data/service)

- persistence: once given a name, the name remains valid
  - today: HTTP redirect + dynamic DNS
- availability: high available in terms of latency & reliability
  - today: endpoint replication + network load balancing (find nearby copy)
- authenticity: data/service from the appropriate source
  - today: secure channels to the source

today's solution

ad-hoc, application-specific, expensive work-around

- today: HTTP redirect + dynamic DNS

- today: endpoint replication + network load balancing (find nearby copy)

- today: secure channels to the source

today's solution

ad-hoc, application-specific, expensive work-around

- to

data-centric applications made unnecessarily hard by the host-to-host Internet architecture

- to arby copy)

- today: secure channels to the source

DONA — data-oriented network architecture

centered around how Internet names are structured and resolved

key idea

- replace DNS name-based resolution with a name-based anycast primitive

DONA — data-oriented network architecture

basic design and usage

broader impacts

DONA basic design

strict separation of concerns between naming and name resolution

DONA basic design

strict separation of concerns between naming and name resolution

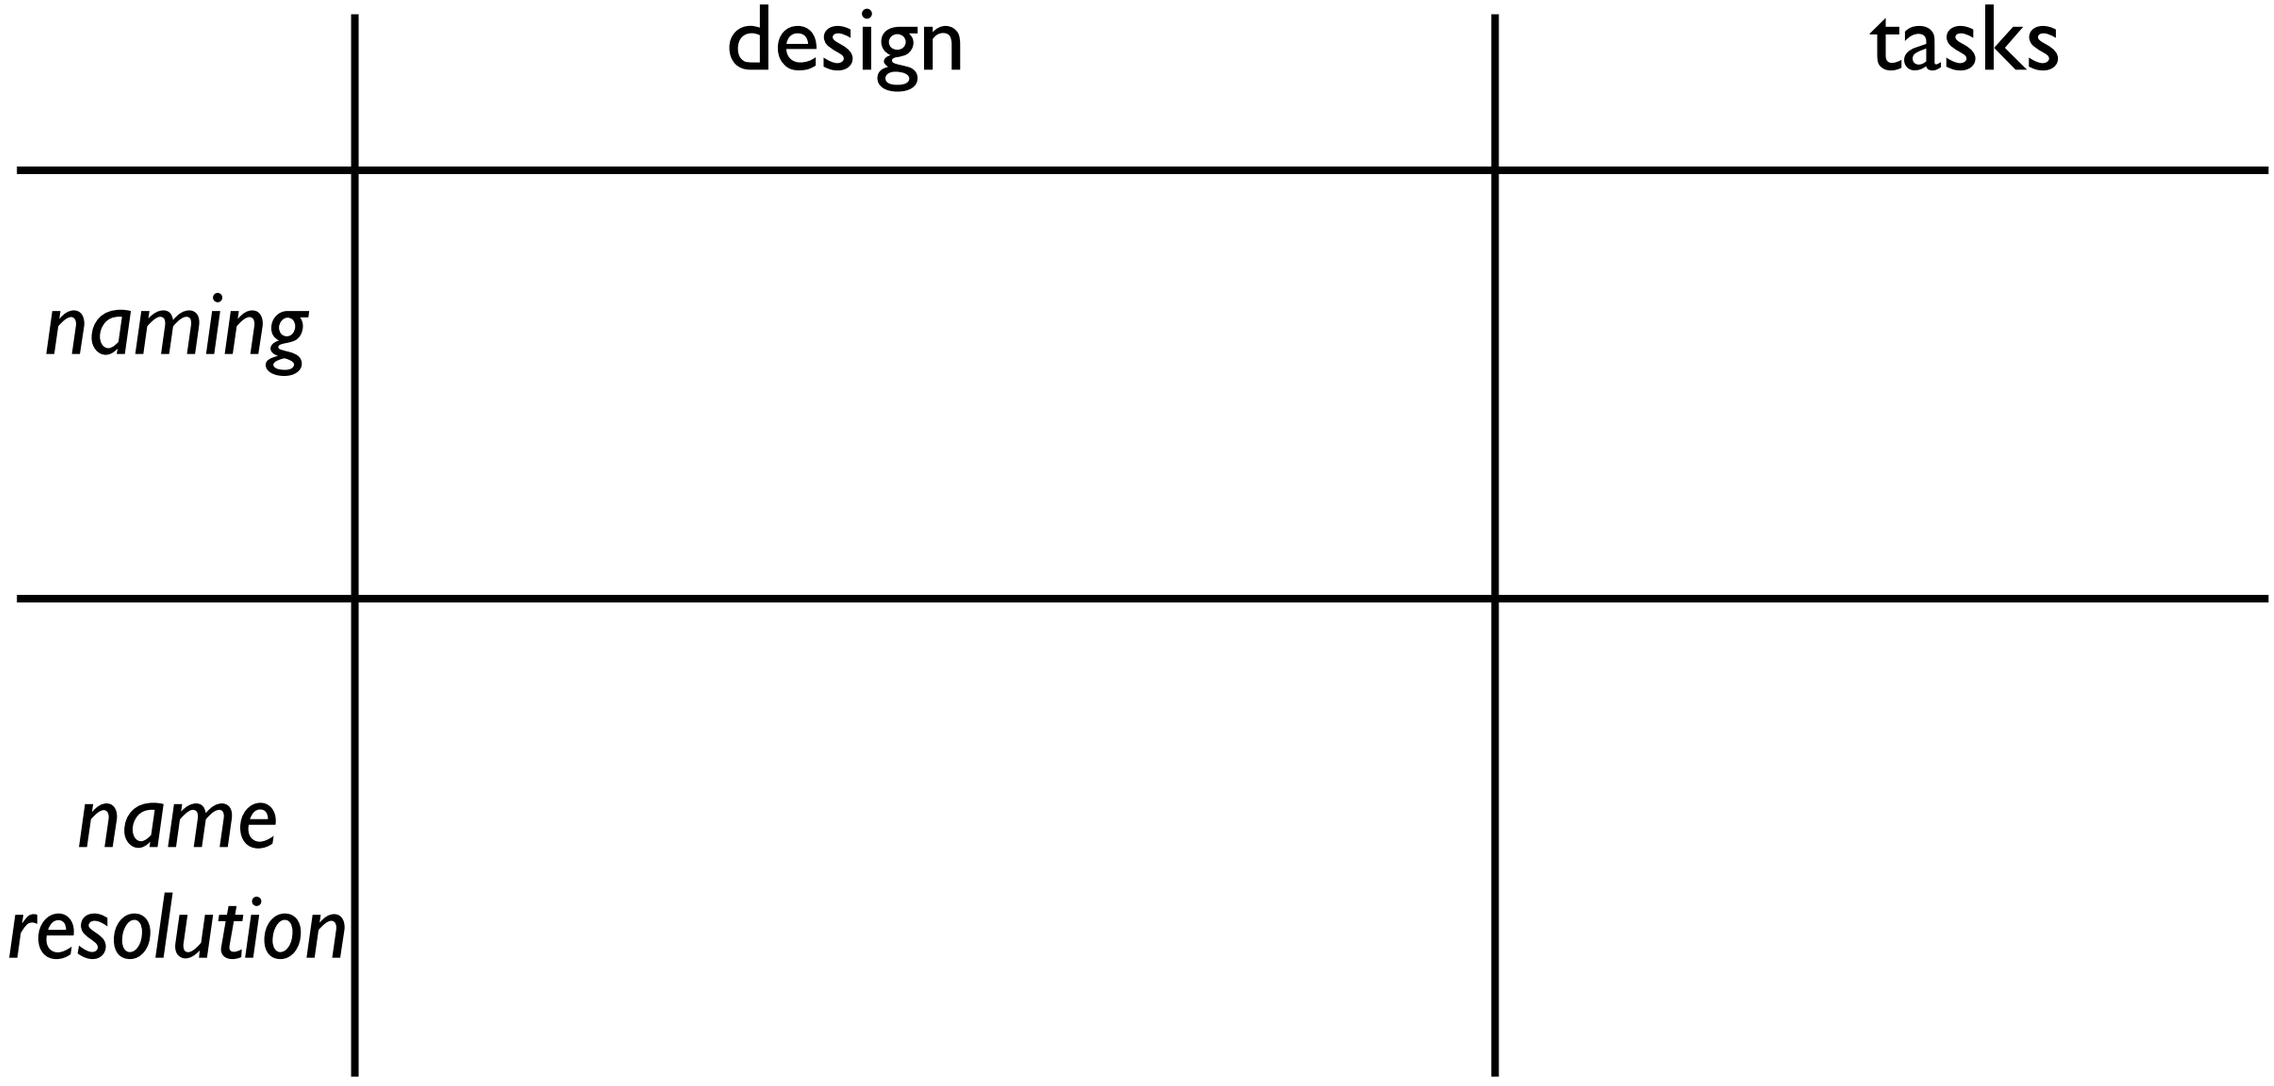

DONA basic design

strict separation of concerns between naming and name resolution

	design	tasks
naming	~~DNS names~~ flat, self-certifying names	persistence authenticity
name resolution		

DONA basic design

strict separation of concerns between naming and name resolution

	design	tasks
naming	~~DNS names~~ flat, self-certifying names	persistence authenticity
name resolution	~~lookup-by-name in a distributed database~~ route-by-name to the nearest copy	availability • guide requests to nearby copies • avoid failed / overloaded servers

naming

P : L

principle:
cryptographic hash of principle's public key

unique label
chosen by the principle

principle P: owns the data/service, only hosts authenticated by P can offer the data/service

naming

P : L

principle: cryptographic hash of principle's public key

unique label chosen by the principle

self-certifying

- client (host) asks data with name P:L, receives <data, public key, signature>
- client verifies data is from P by checking if the public key hashes to P, if the key generates the signature

naming

P : L

principle:
cryptographic hash of principle's public key

unique label
chosen by the principle

challenge

- how to resolve to the appropriate location

name resolution

DNS approach

- lookup-by-name in a distributed database
- given a name, returns the location (IP address) of a nearby copy

DONA approach

- route-by-name
- rely on a new class of network entities — resolution handler (RH) and two basic primitives: FIND(P:L), REGISTER(P:L)

name resolution primitives

FIND(P:L), REGISTER(P:L)

FIND packet locates the object name P:L

REGISTER message sets up the state necessary for the RHs to route FINDSs effectively

resolution handlers (RHs)

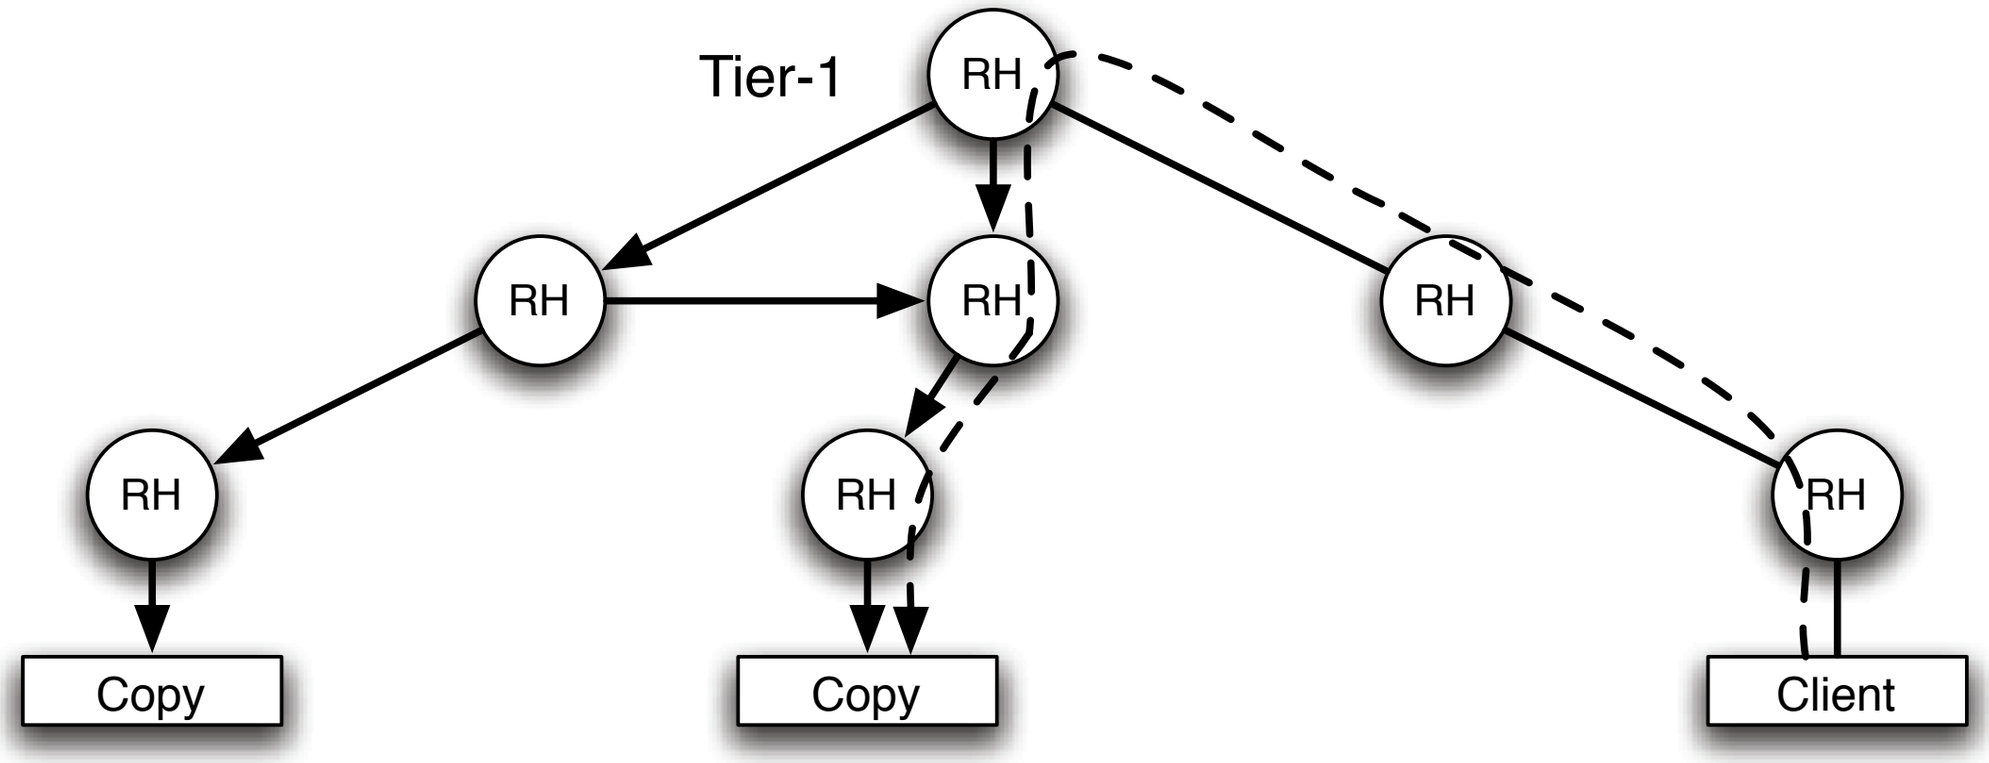

each domain has a (local) RH

- RHx: resolution handler associated with a domain X
- RHx is the provider/peer/customer of RHy if X is the provider/peer/customer of Y

resolution handlers (RHs)

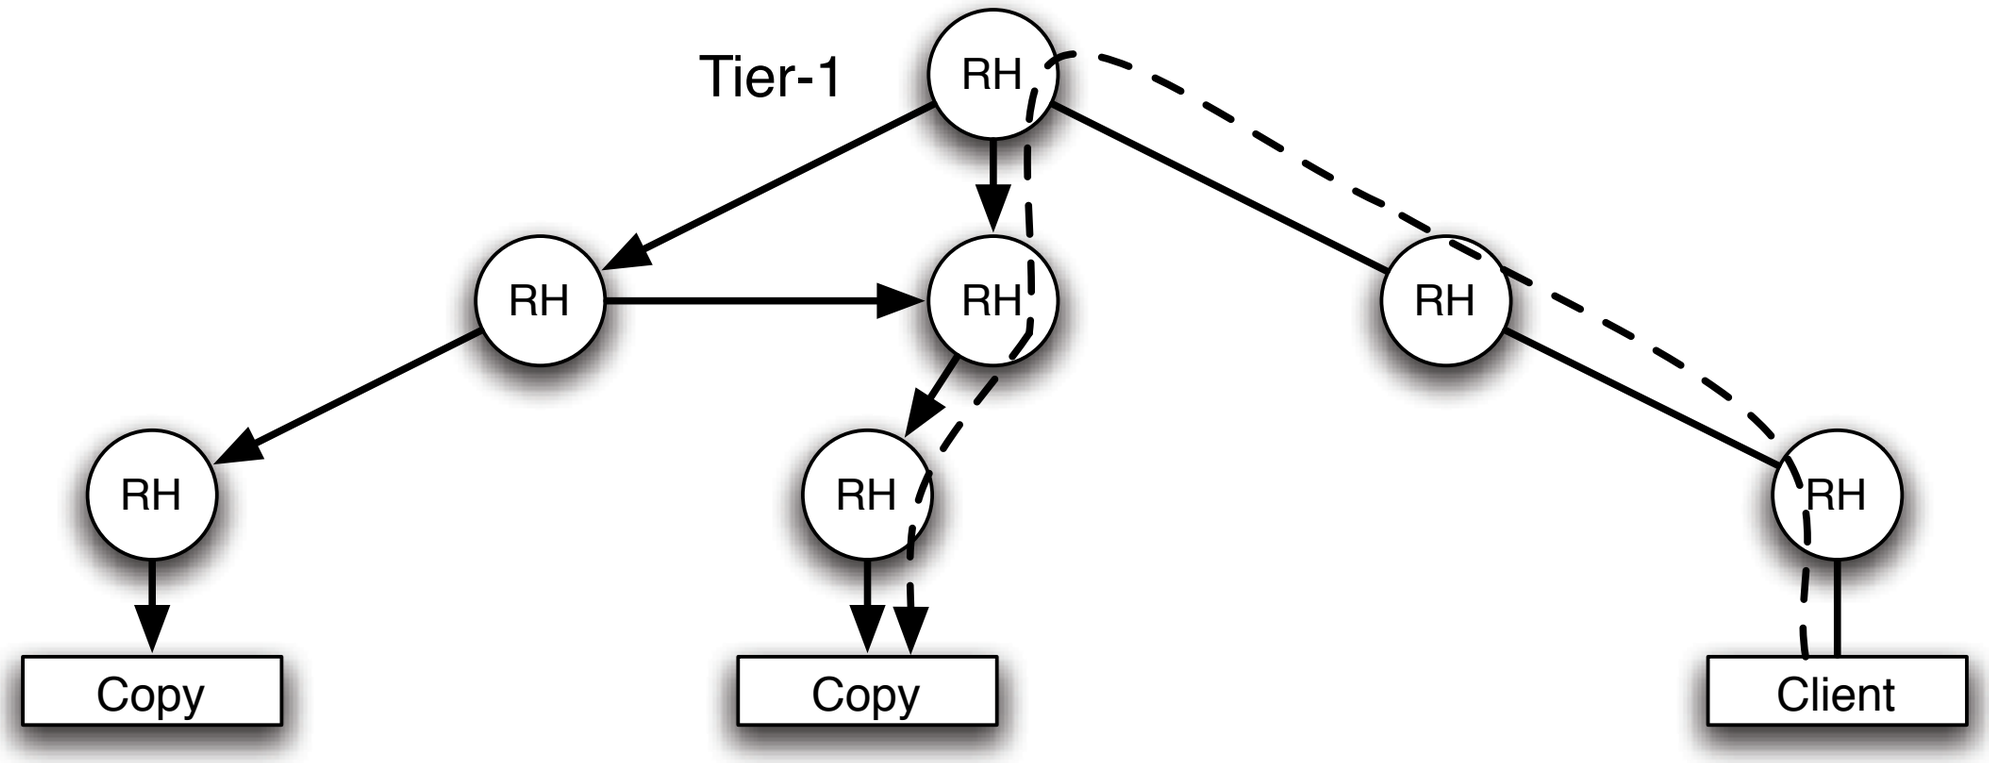

each domain maintains a registration table

- maps a name (P:L) to the next-hop RH and the distance to the data copy

RH primitive — FIND(P:L)

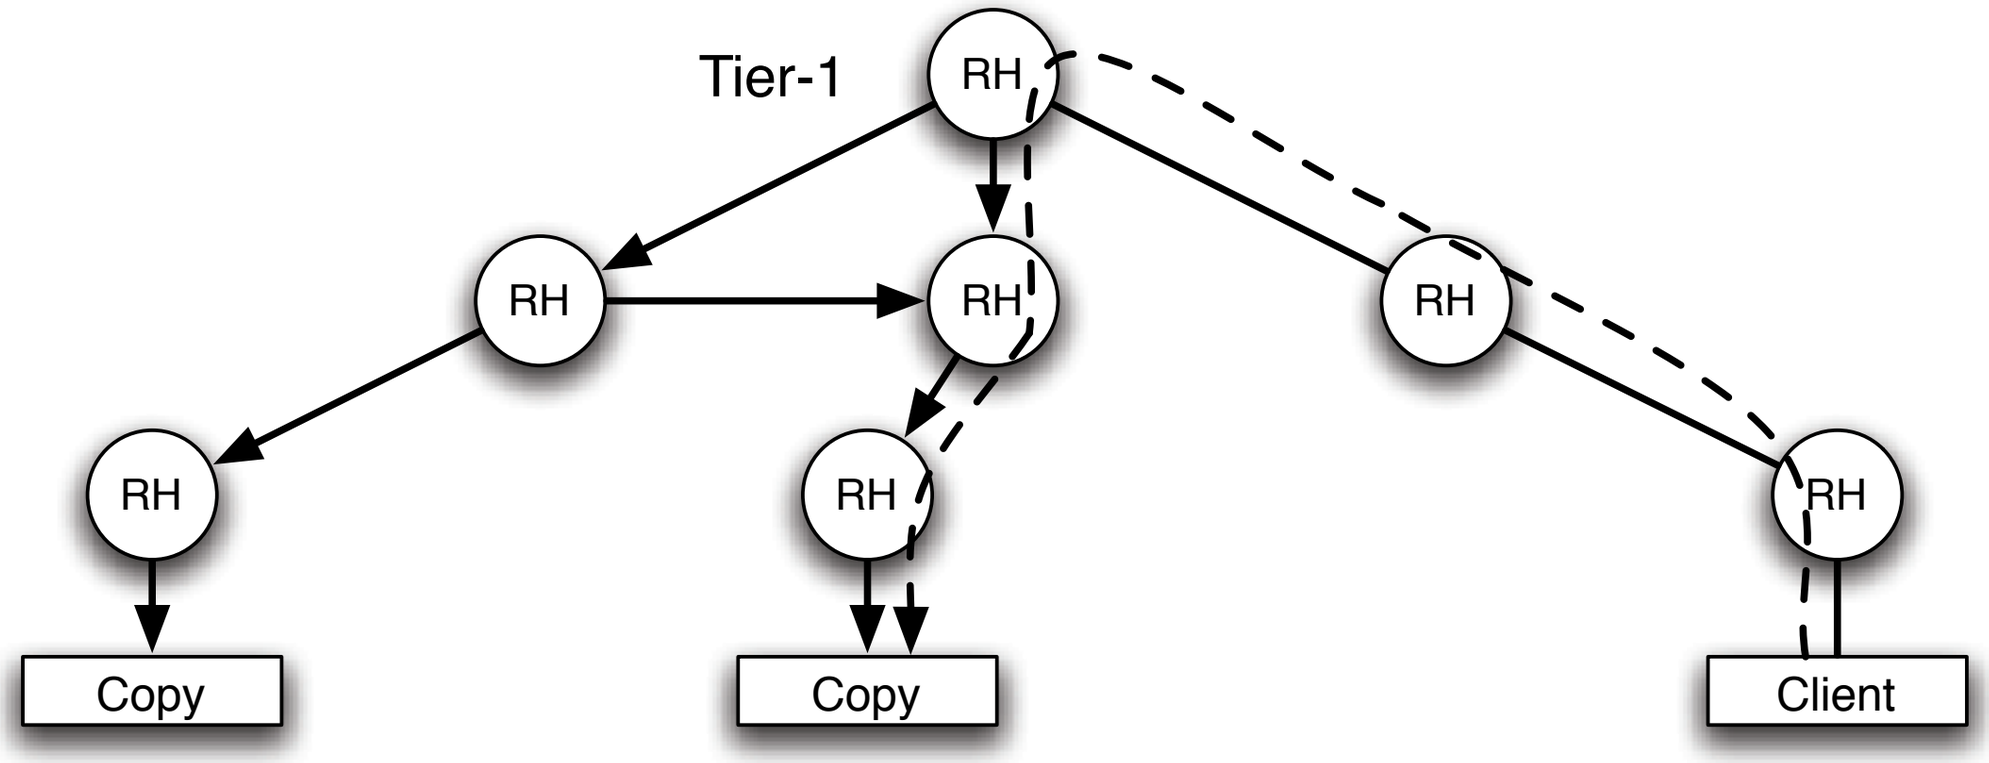

RHs route client-issued FIND to a nearby copy

- is there an entry matches (P:L) in the local registration table?
- YES: route to the next-hop RH
- NO: forwards the FIND message up the RH hierarchy
  - in the hope of finding an entry

RH primitive — REIGSTER

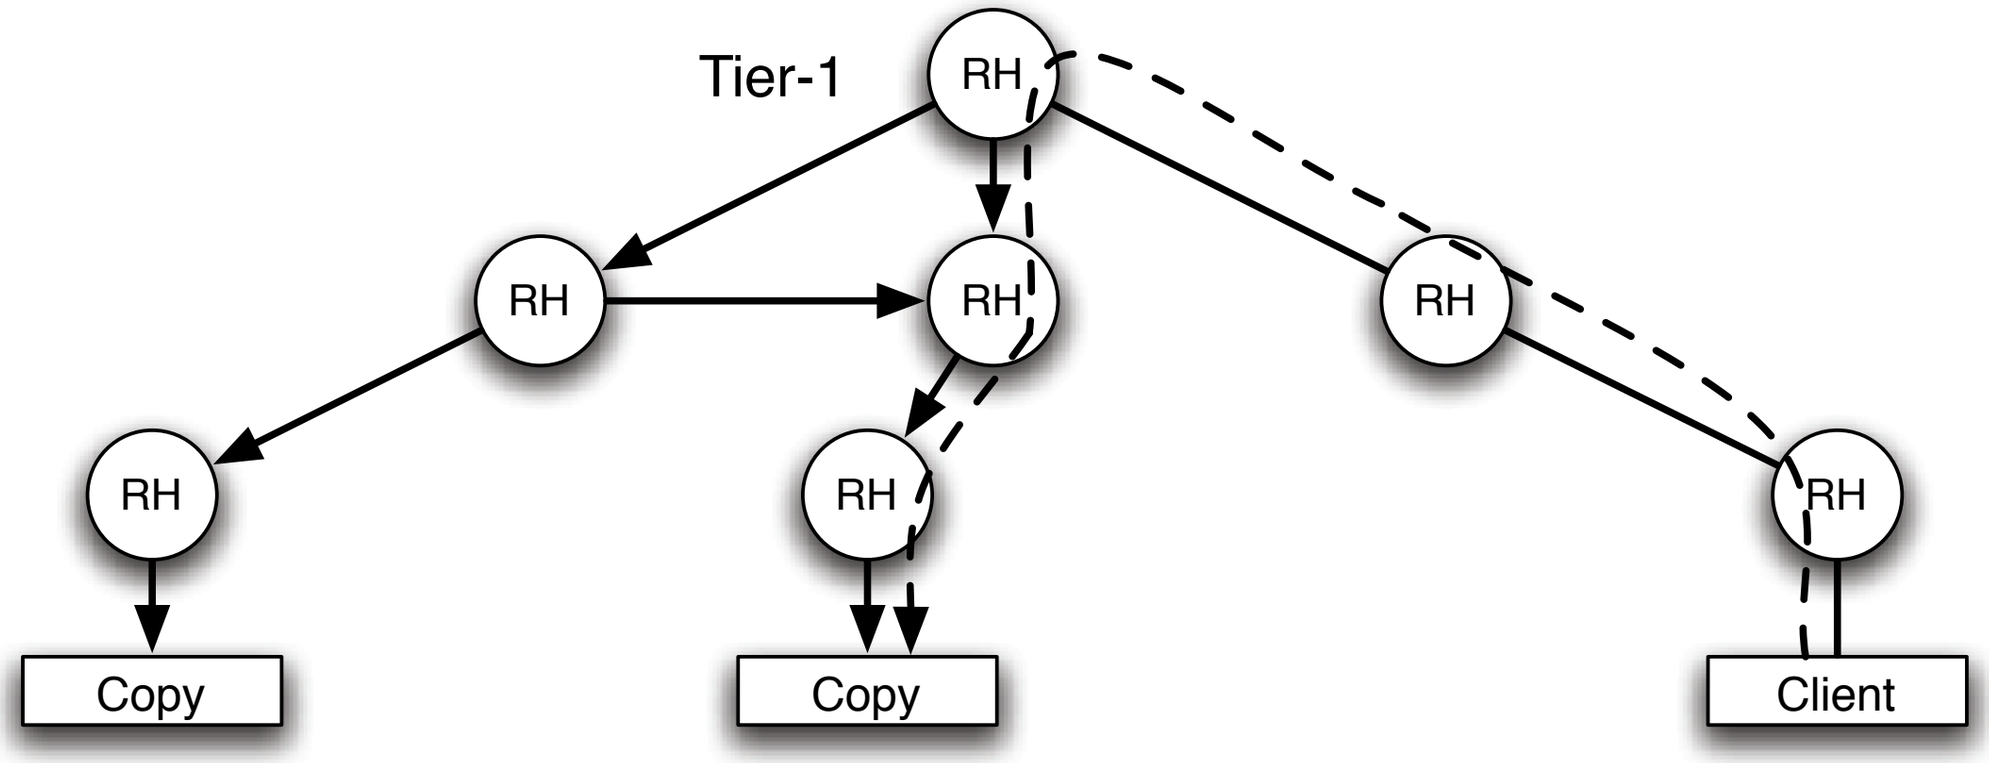

RHs receive a REGISTER message from a child (customer)

- if no entry exists, no nearby copies
  - updates its own registration table
  - forwards the REGISTER message to its parents and peers

using DONA

DONA is essentially a name-based any cast service: a FIND(P:L) request is routed to nearby RH at which the name P:L is registered

- service selection among several possible servers
  - DONA route FIND(P:L) to the closest server

using DONA

DONA is essentially a name-based any cast service: a FIND(P:L) request is routed to nearby RH at which the name P:L is registered

- service selection among several possible servers
  - DONA route FIND(P:L) to the closest server
- mobility
  - a roaming host first unregisters from old location
  - re-registers at new location
  - subsequent FINDs routed to the new location

using DONA

DONA is essentially a name-based any cast service: a FIND(P:L) request is routed to nearby RH at which the name P:L is registered

- service selection among several possible servers
  - DONA route FIND(P:L) to the closest server
- mobility
  - a roaming host first unregisters from old location
  - re-registers at new location
  - subsequent FINDs routed to the new location
- multihoming
  - multi-homed hosts register with each local RH
  - multi homed domain forwards REGISTERs to each provider
  - subsequent data-connections can make use of multiple paths

broader impacts

middlebox — another issue where historical design is at adds with current usage

- Internet architecture follows end-to-end argument
  - a purely transparent carrier of packets
- intermediary routers playing no role other than forwarding packets

broader impacts

middlebox — another issue where historical design is at adds with current usage

- but on-path middle boxes violate end-to-end argument

broader impacts

middlebox — another issue where historical design is at adds with current usage

- but on-path middle boxes violate end-to-end argument

unmet needs of middleboxes due to commercial, security pressures

extend DONA to support middlebox

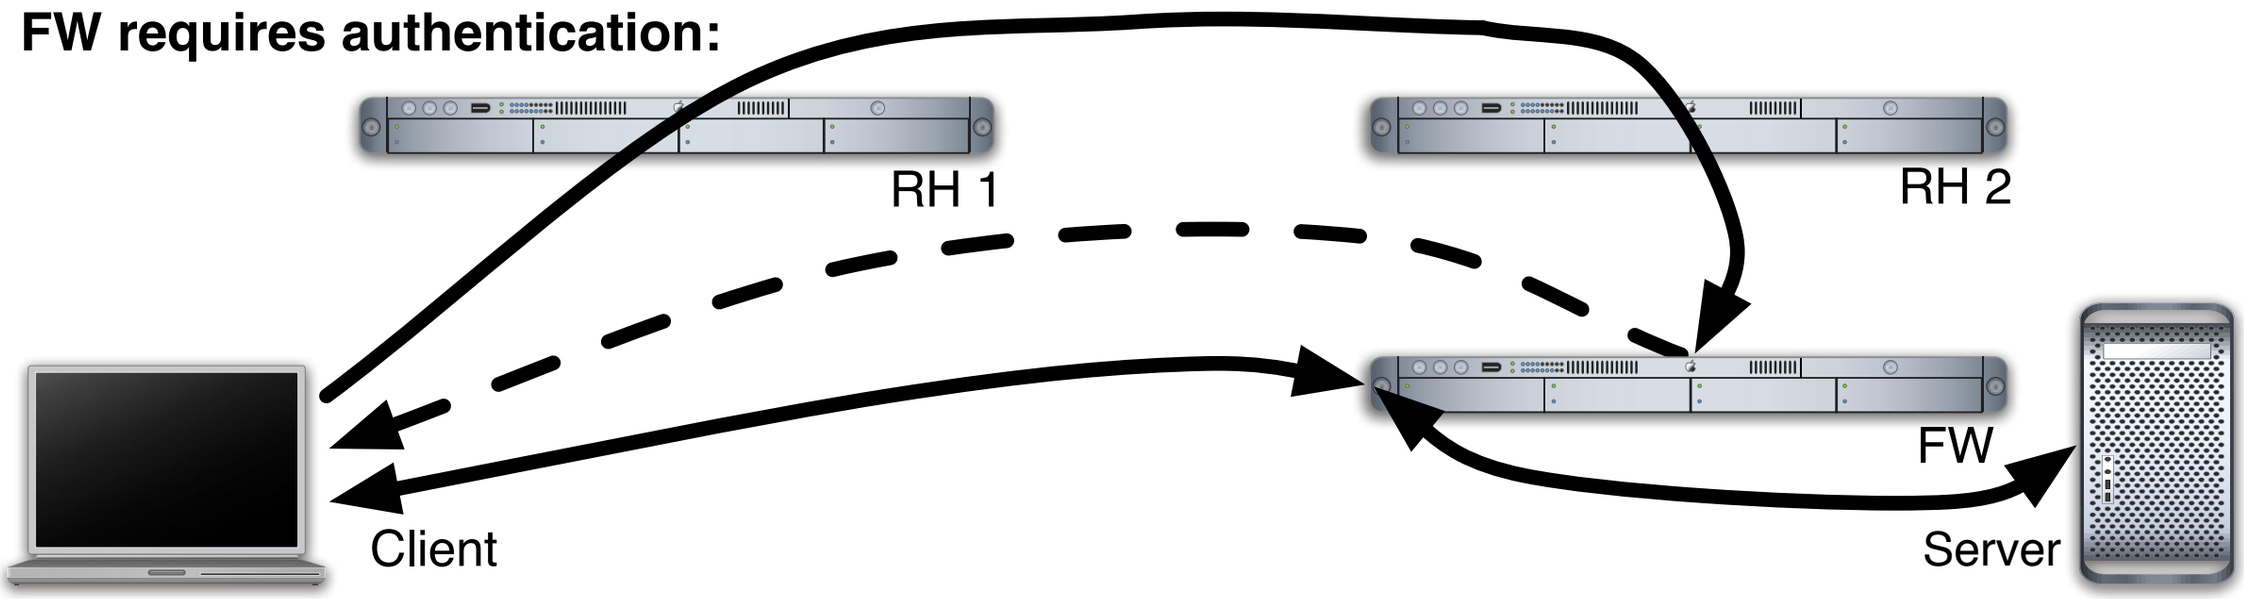

DONA treats stakeholders along the path as relevant Internet actors

extend DONA to support middlebox

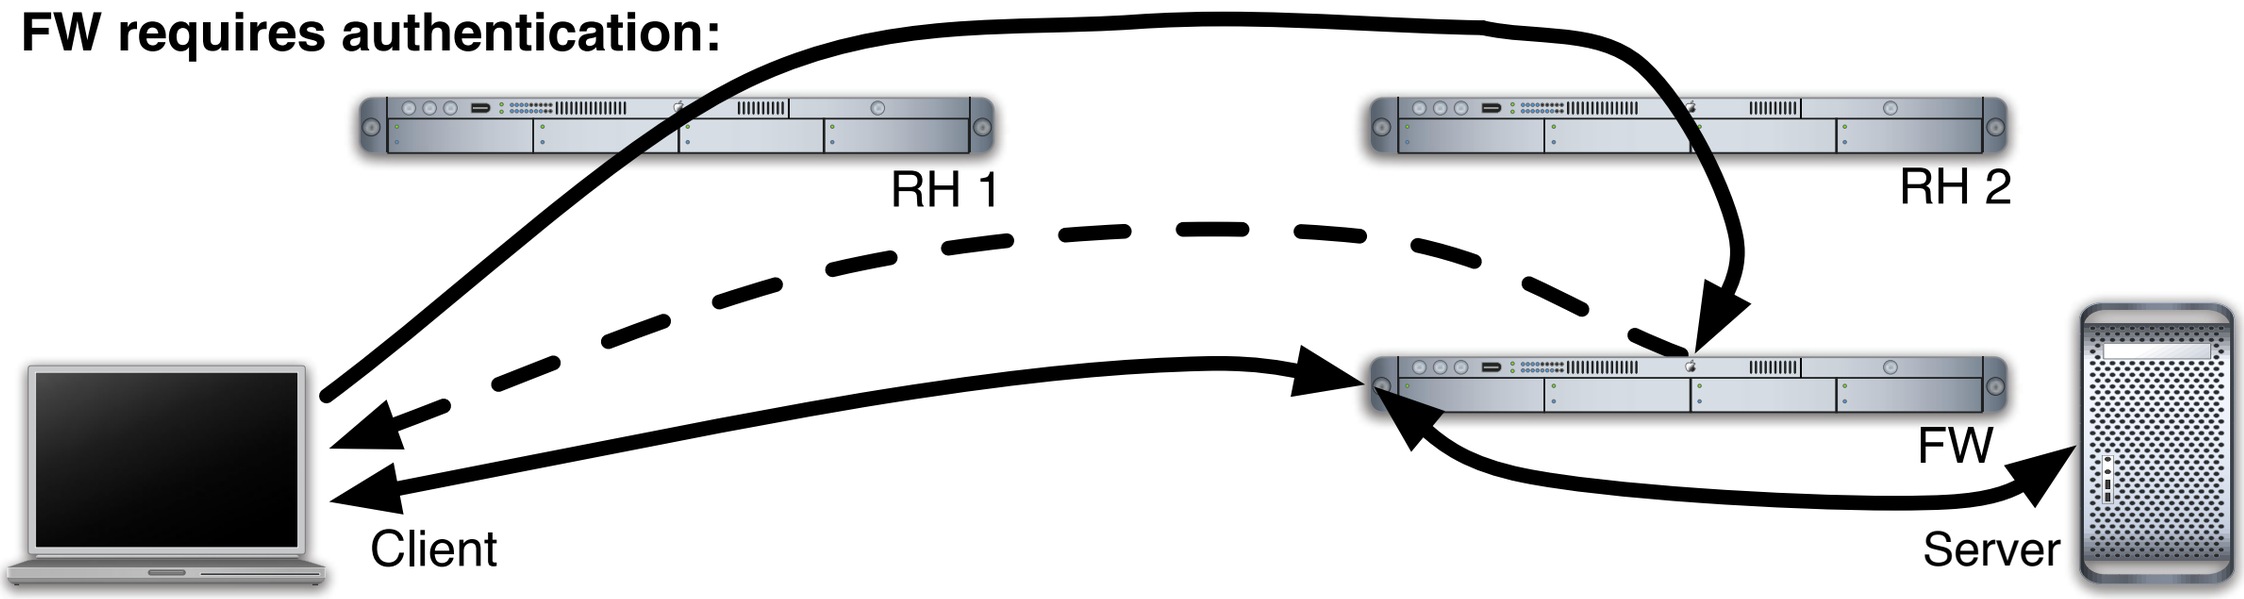

FIND goes from client RH1 to RH2

- RH2 forwards it to the firewall FW
- FW asks client for authentication
- client responds and FW verifies
- FW becomes proxy for client-server communication