

# Building a Coreless Internet Without Ripping Out the Core

**Geoffrey Goodell**

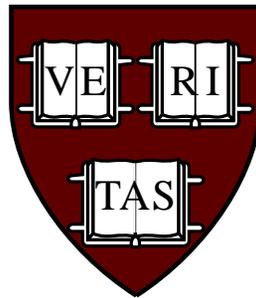
(goodell@eecs.harvard.edu)

**Scott Bradner**

(sob@harvard.edu)

**Mema Roussopoulos**

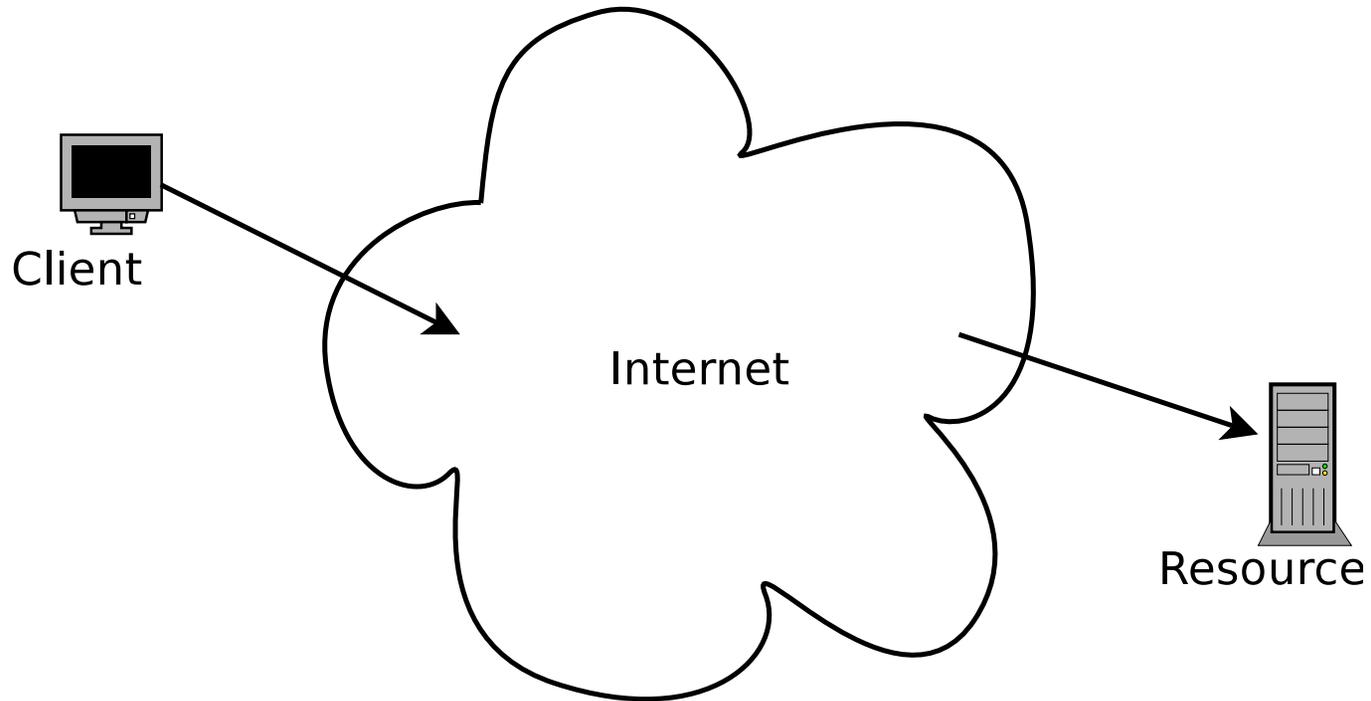
(mema@eecs.harvard.edu)



Harvard University

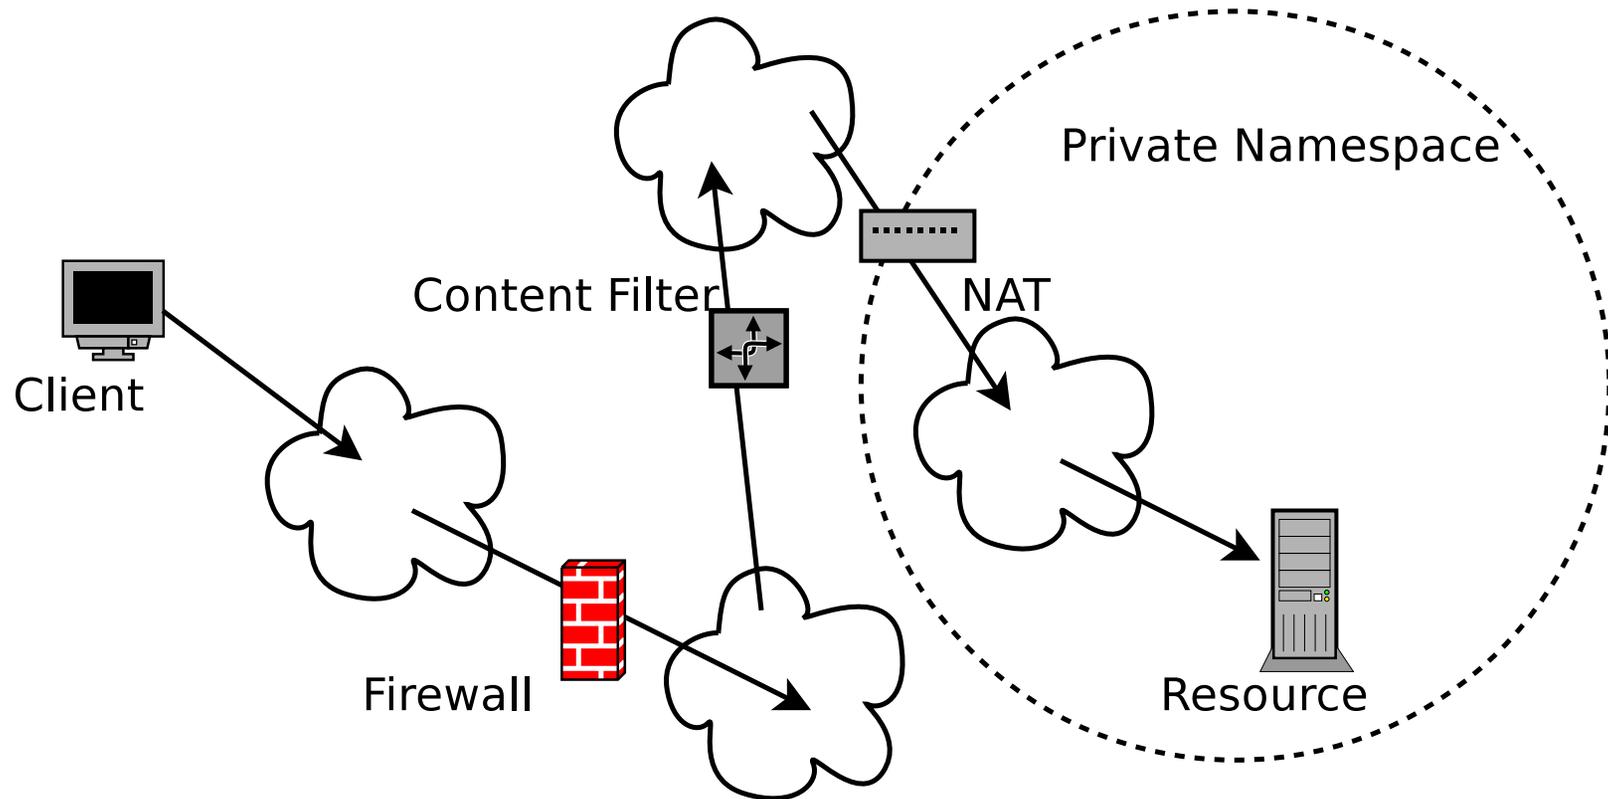
HOTNETS-IV: 14 November 2005

## The Dream



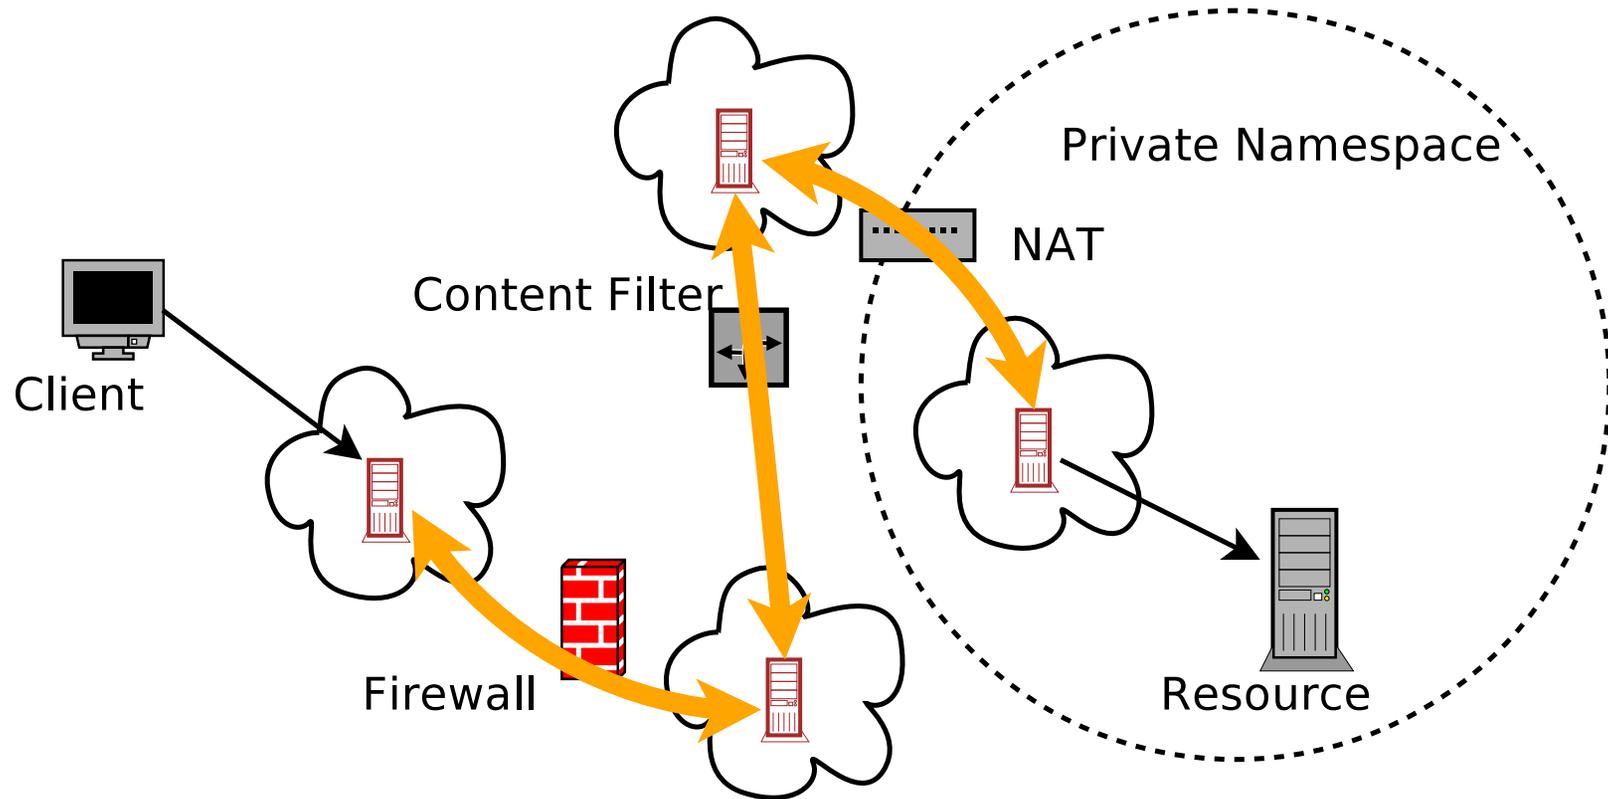
- All clients have access to all resources.

## The Reality



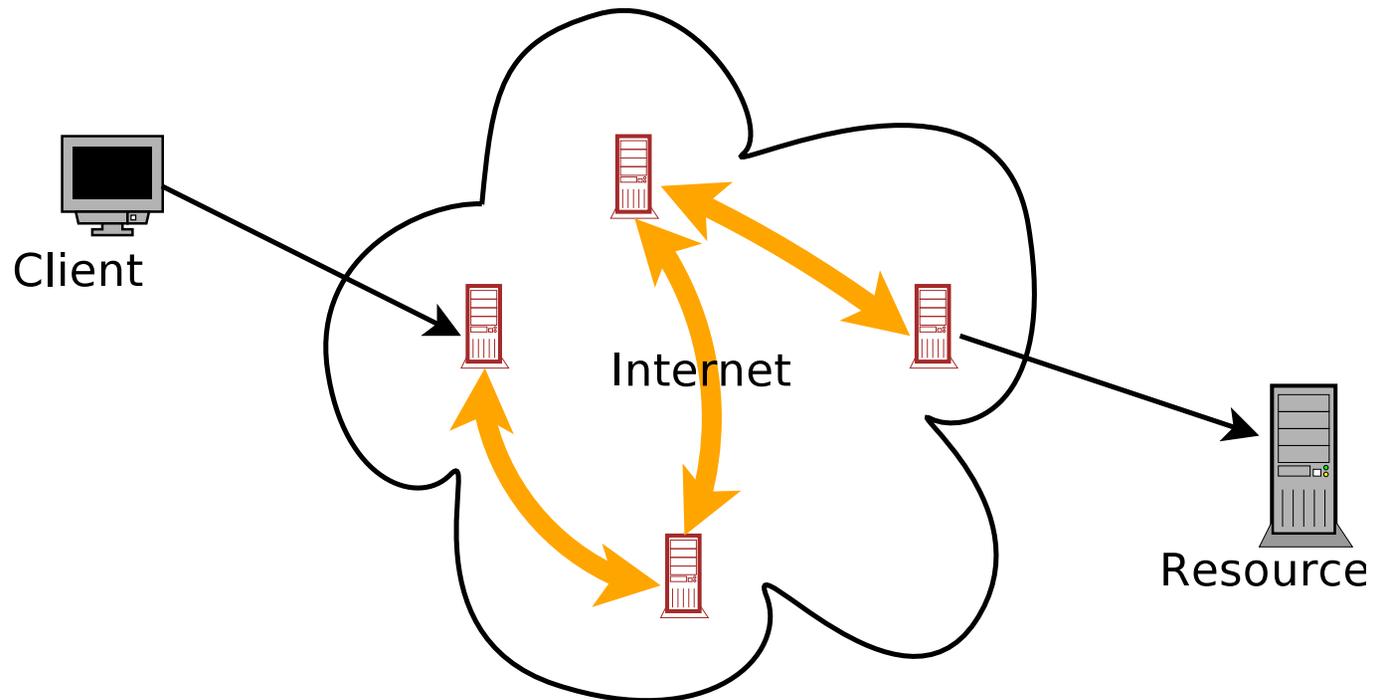
- Access to network resources is a function of location within the topology.

# Blossom Approach



- A peer-to-peer overlay network for sharing **perspectives**.

# Blossom Vision



- A **Coreless** Internet: the meaning of names and addresses is a function of their context.

“The IP suite imposes a single networking model and addressing scheme over the many different underlying network types it supports... Although helpful for a number of years, while the network was undergoing its initial ‘big bang’ phase, the approximations in the model **are now becoming too out-of-step with reality.**”

—Jon Crowcroft et al., *Plutarch*, 2003

“Without consensus, some experts say that countries might move ahead with **setting up their own domain name system**, or DNS, as a way of bypassing Ican.

“The United States argues that **a single addressing system is what makes the Internet so powerful**, and moves to set up multiple Internets would be in no one’s interest.”

—*Tom Wright, “EU Tries to Unblock Internet Impasse.”*

*The New York Times, 30 September 2005*

## Some Types of Fragmentation

- **Firewalls** (Security)
- **NAT**
  - Network Isolation (Management)
  - Perception of Address Scarcity (Business Model)
- **Others**: Content Filtering, Multiple Namespaces (Political), Anycast (Efficiency), etc.

## The Inherent Conflict

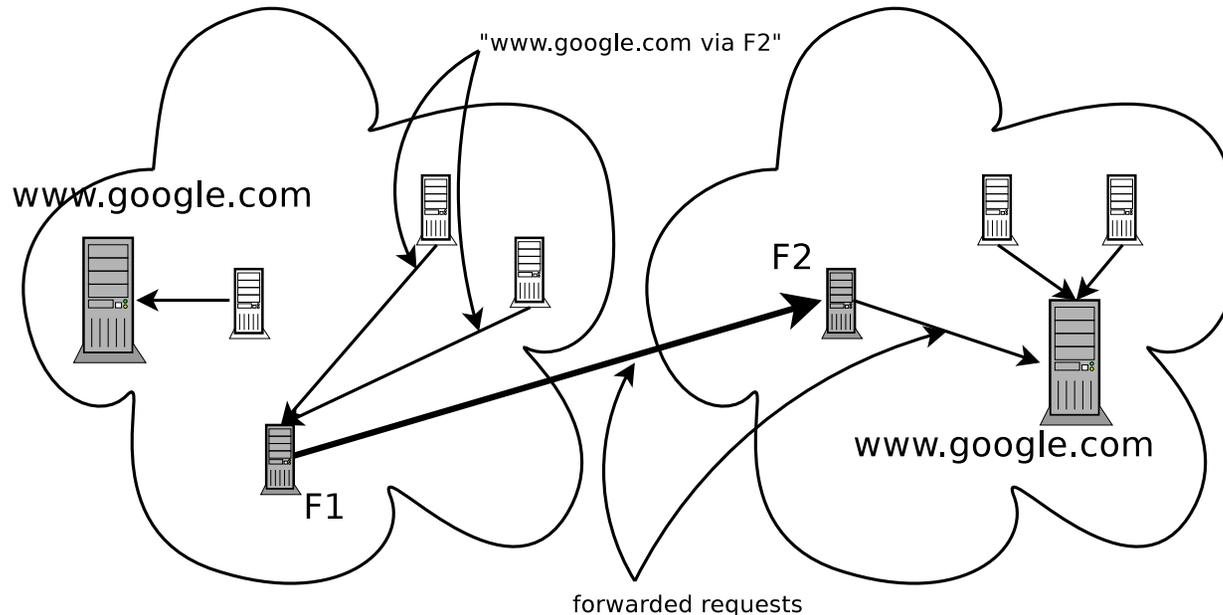
- **Homogenization** promotes universal access and consistency
- **Fragmentation** provides distributed management and location specificity
- Both are **essential** to the success of the Internet

## Blossom Central Principles

- **[1]** Eliminate access restrictions derived from how an observer is connected to the network.
- **[2]** Allow an observer to select a perspective from which to view and access network resources.

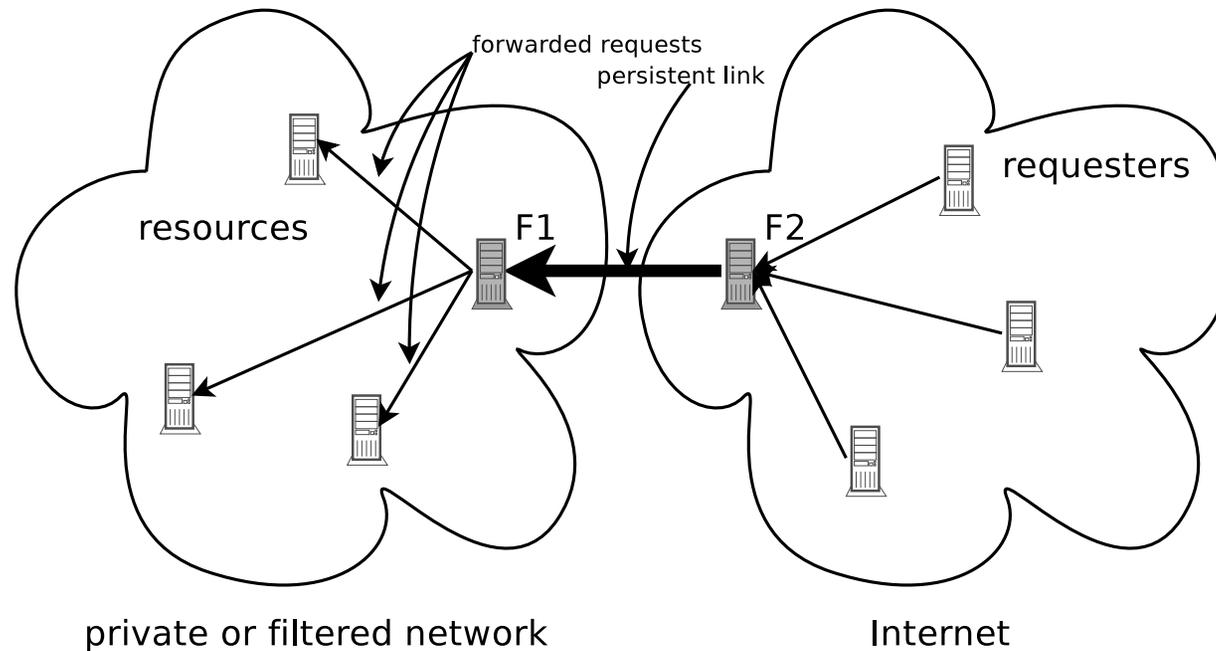
# What We Gain

## What We Gain [1]: Locality



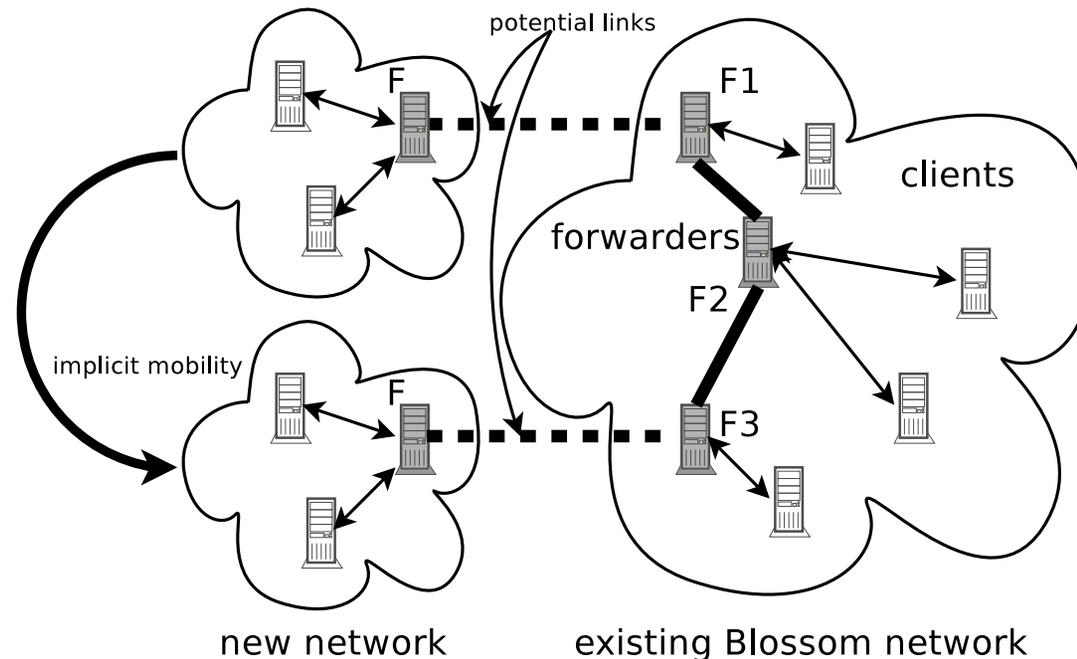
Multiple services with the same name may coexist within different local namespaces (meaningful names within a local space).

## What We Gain [2]: Universal Access



If A can talk to B and B can talk to C, then B can talk to C on behalf of A (circumvent technical barriers).

## What We Gain [3]: Distributed Management

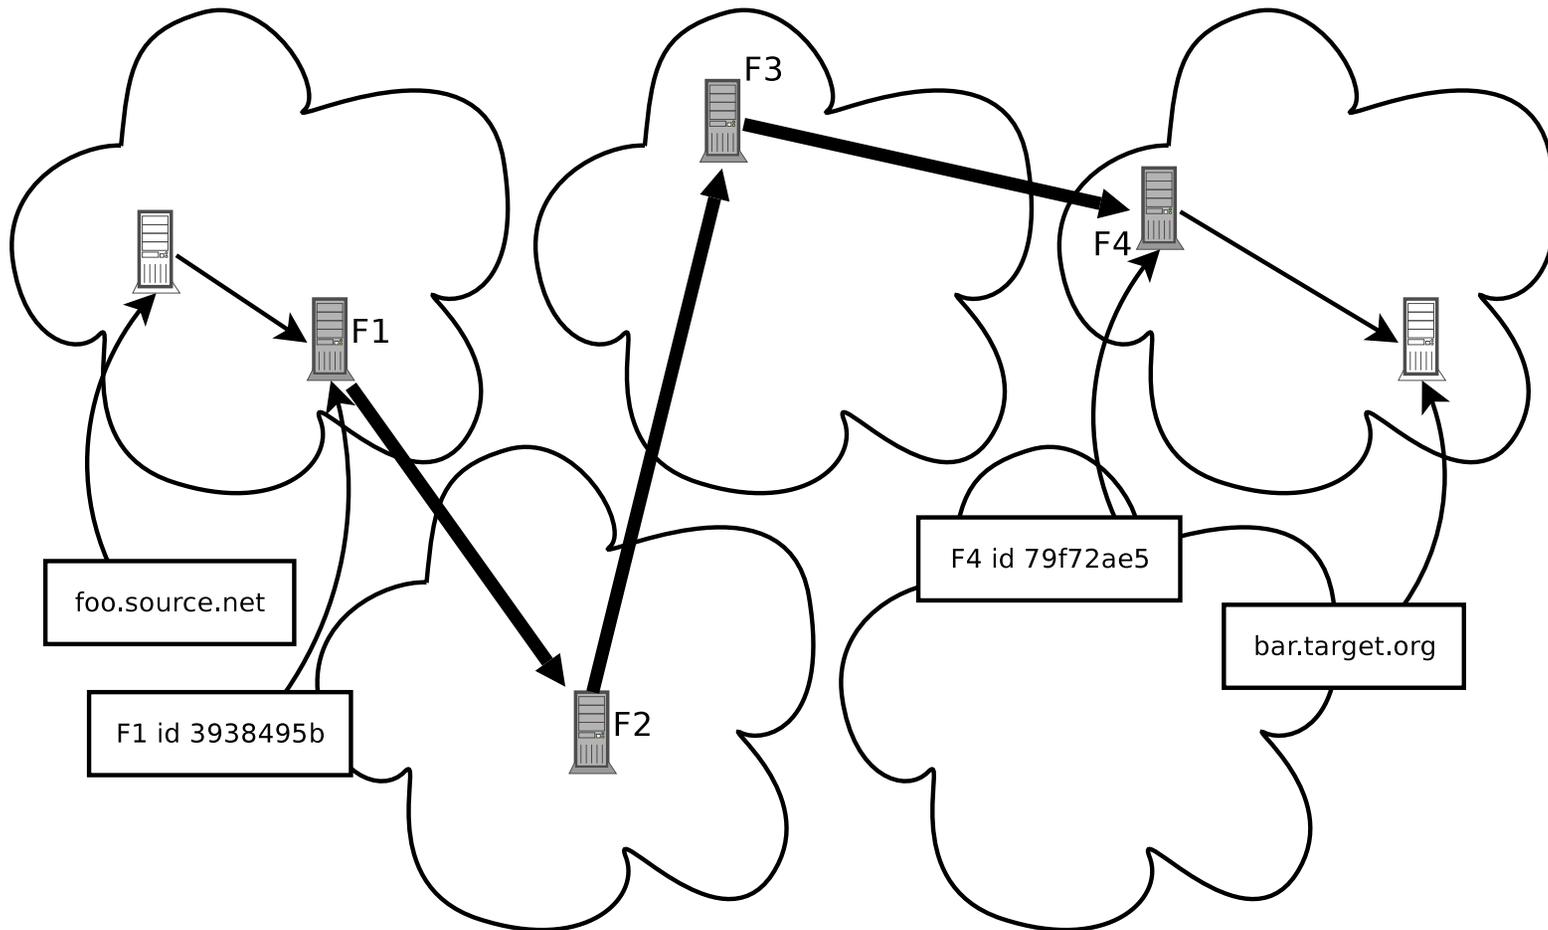


Adding a network and its abundance of resources to the system need not require specific allocation of names, addresses, or routing from centralized authorities.

## What We Gain [4]: Deployability

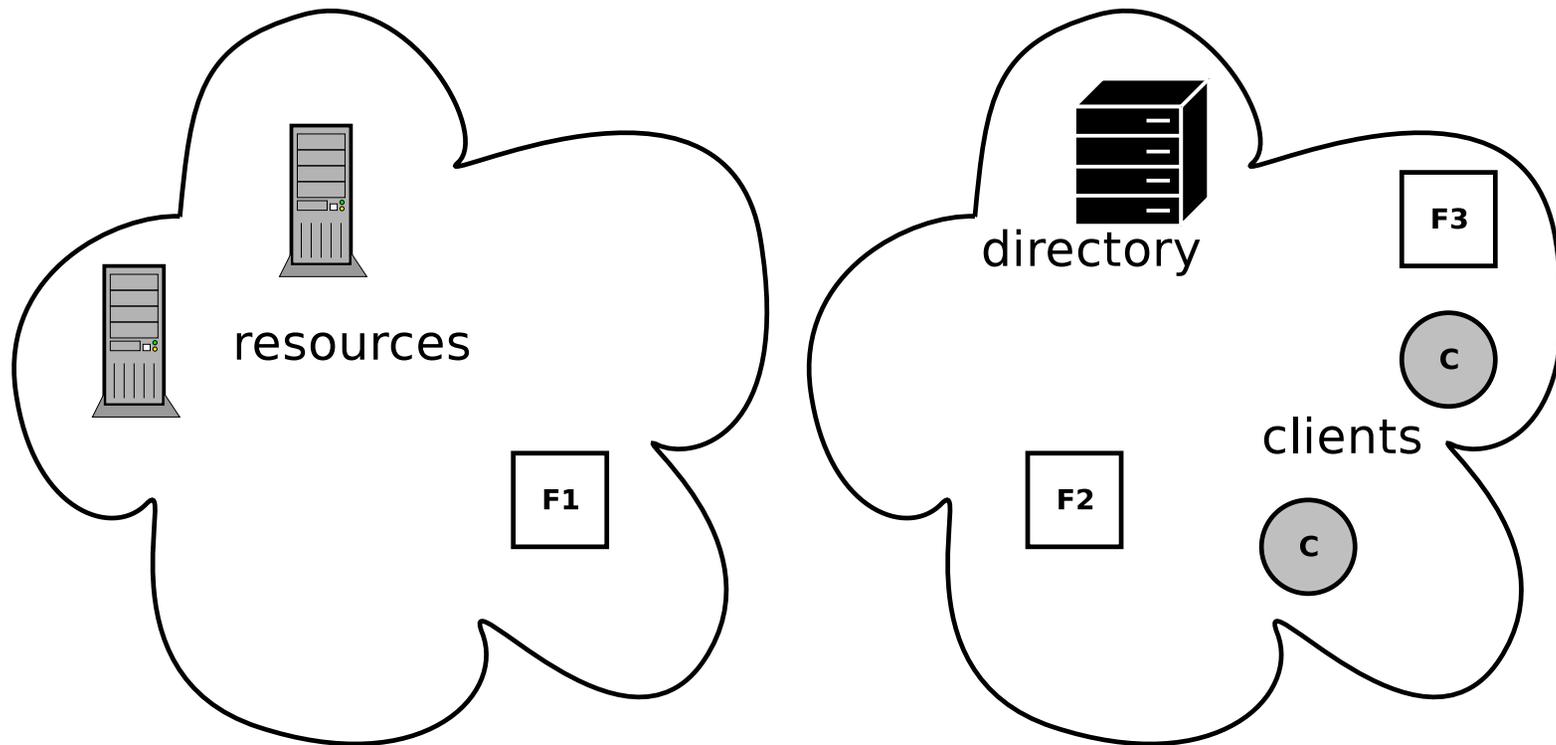
- Must function in the absence of explicit participation of service providers:
  - Network access providers (ISPs)
  - Content providers (legacy Internet services)
- Must provide substantial benefit in the absence of widespread adoption
- Must require minimal (preferably zero) modification to OS or application software

## Accessing a Resource

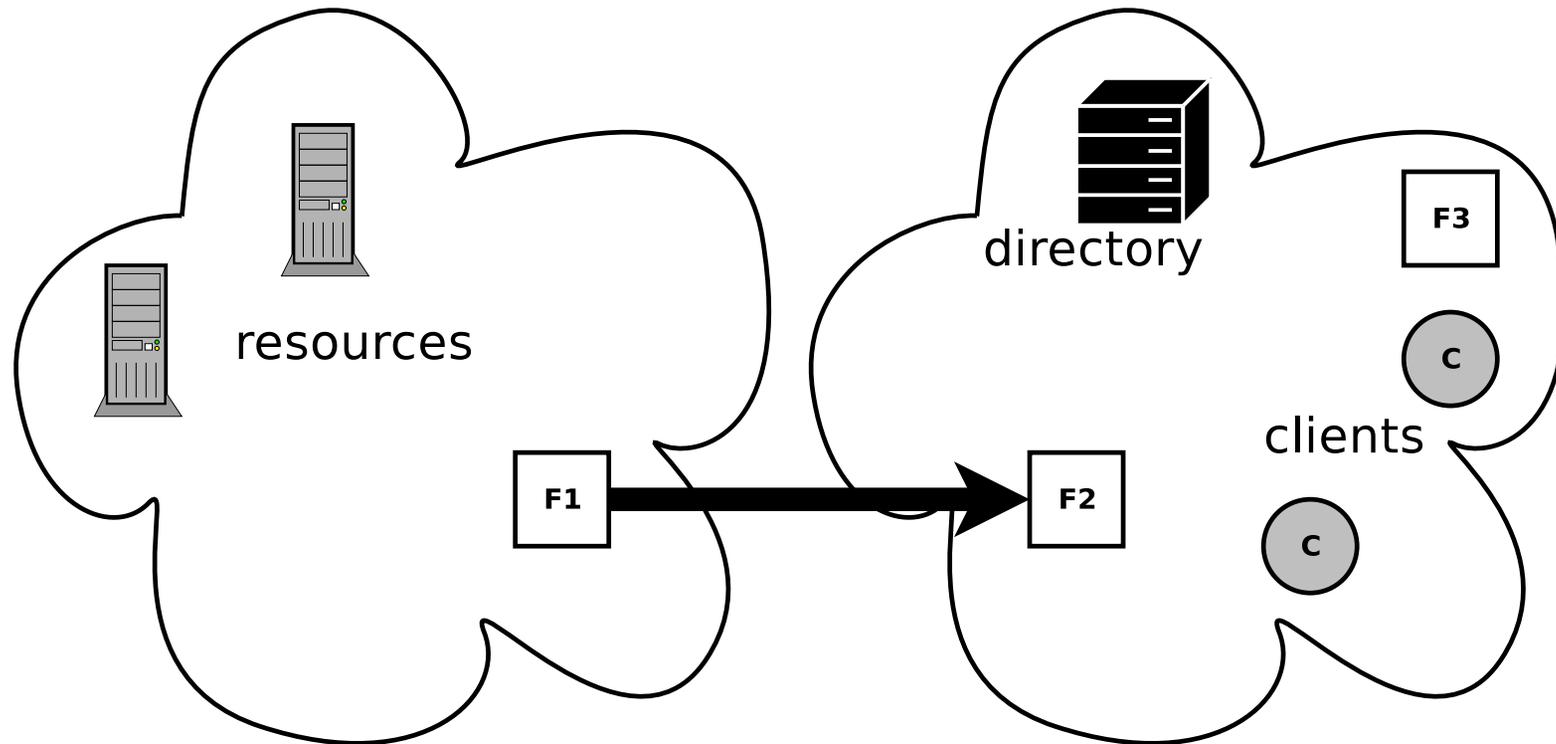


- Our implementation uses **Tor** to build circuits.

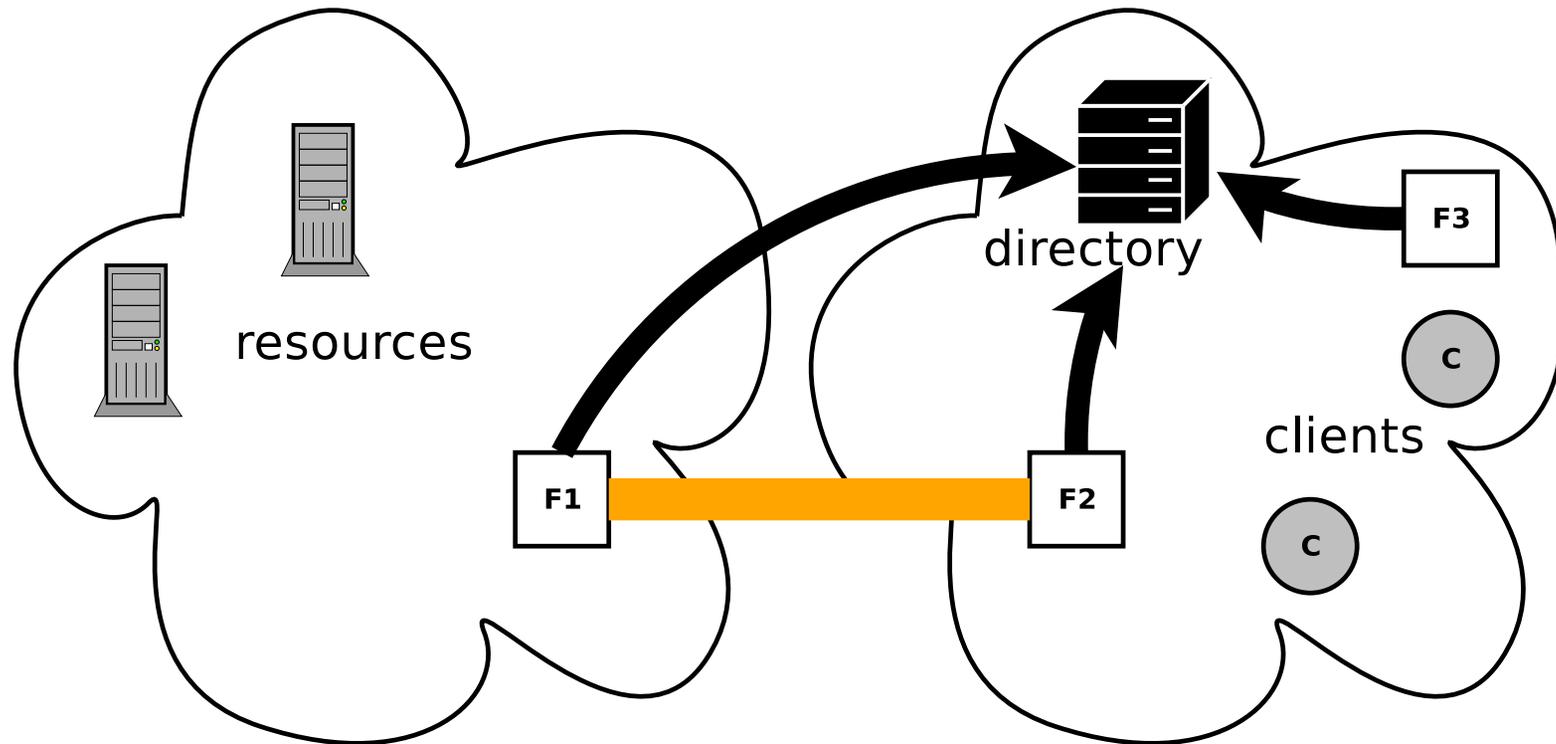
# Blossom Setup



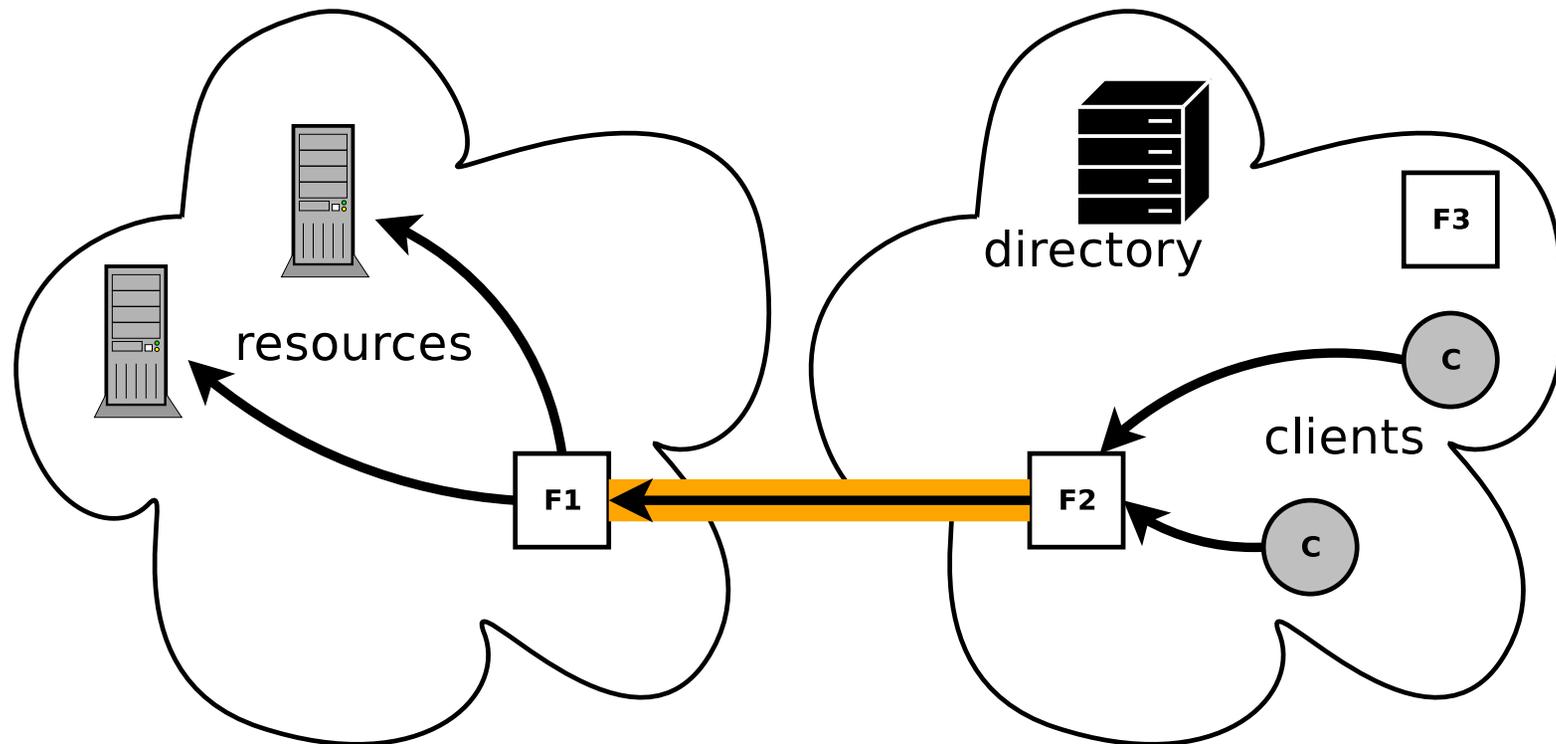
## Establish persistent connection



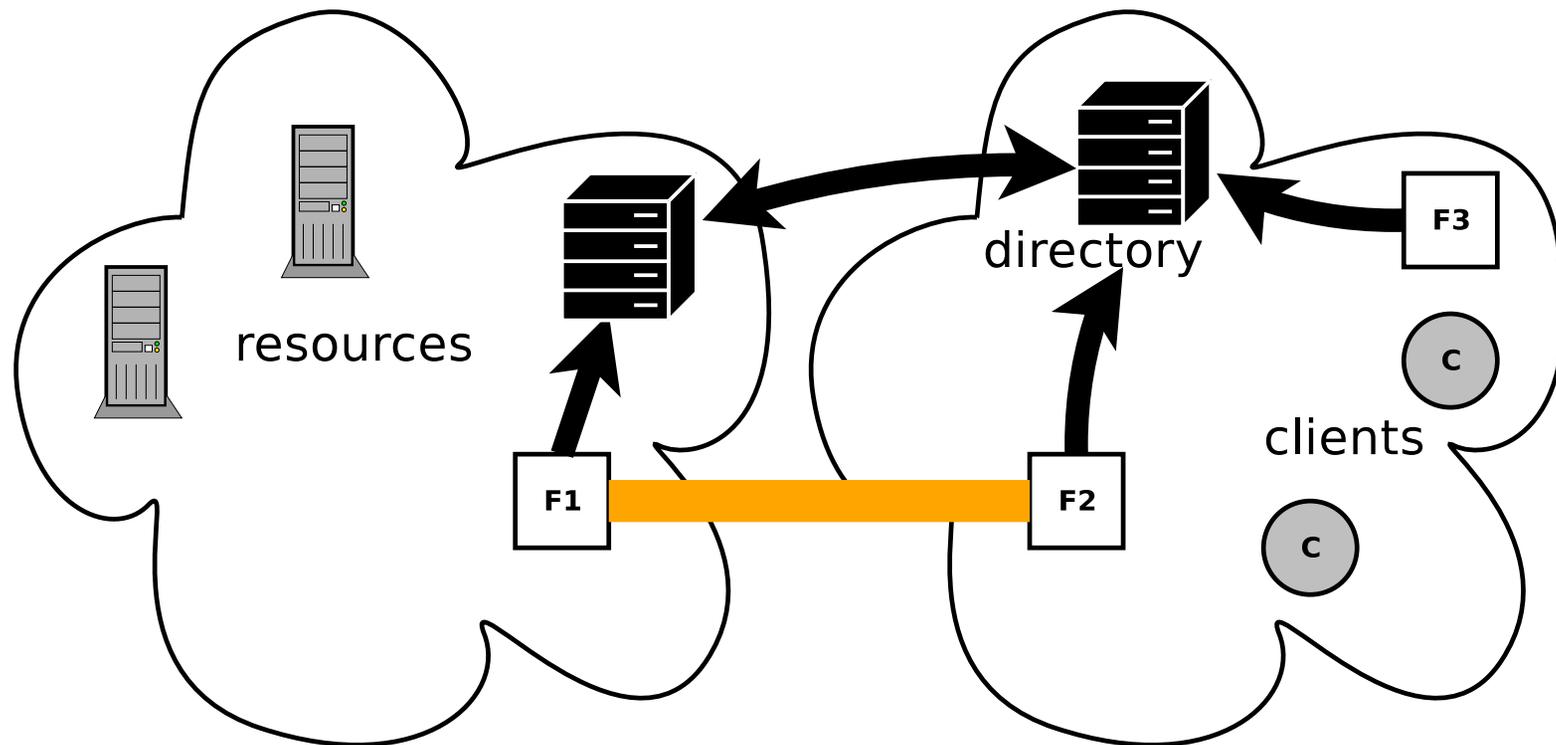
## Publish to directory server



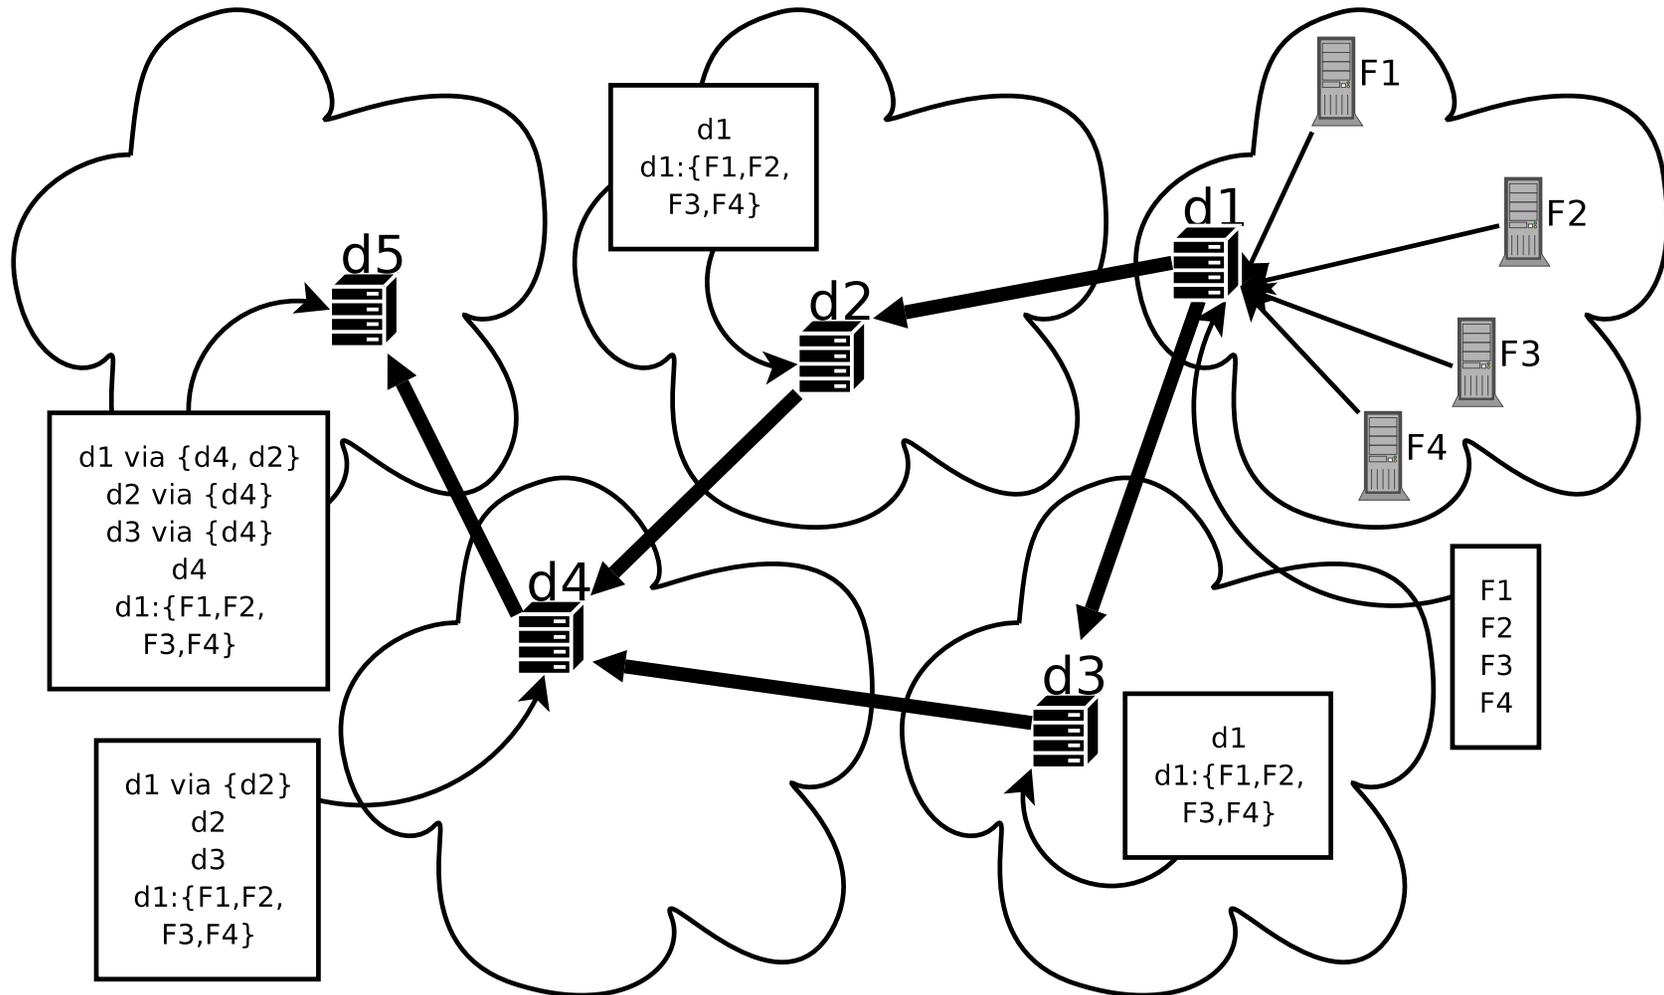
## Clients can now access Resources



# Coreless Approach: LOCAL directory servers



# Directory Propagation (Control Plane)



## Design Tradeoffs (What We Lose)

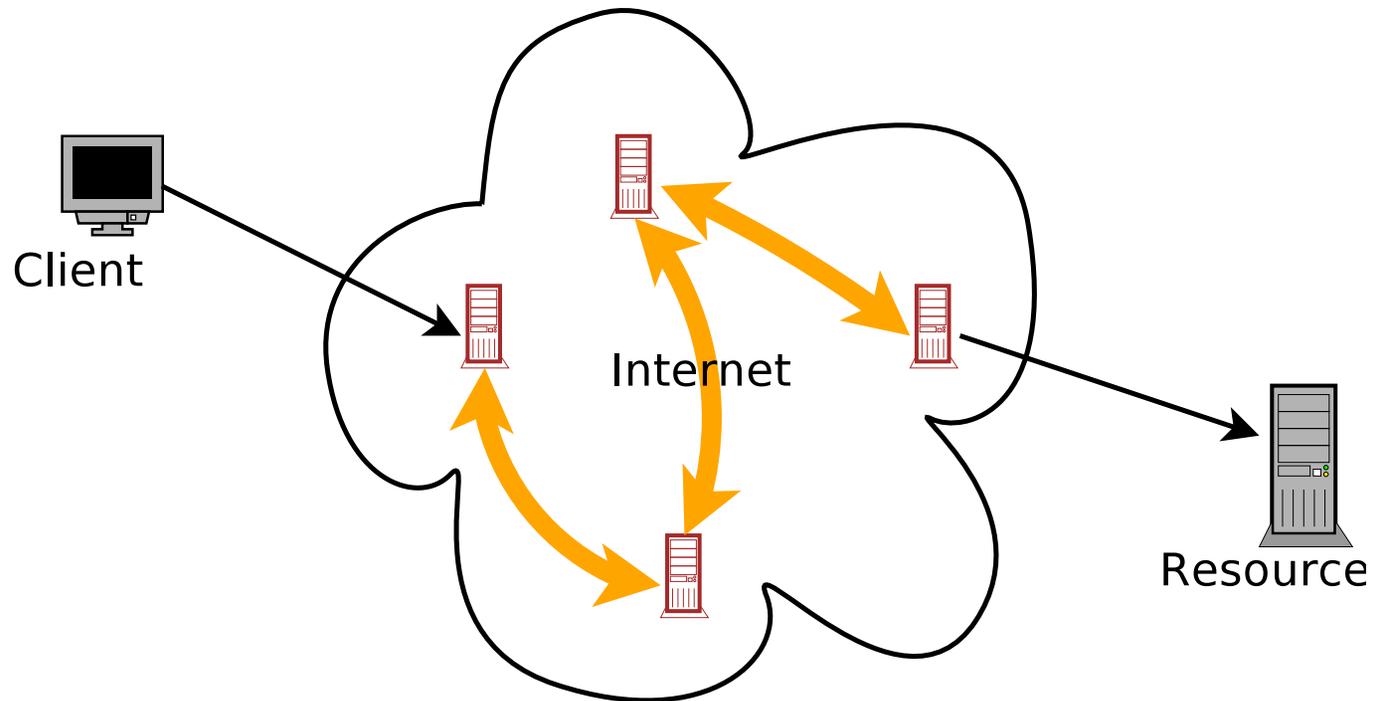
- New **Namespace** Constraints.
  - Do we need globally unique identifiers?
  - Do we need a new DNS?
- New **Discovery** Constraints.
  - Request forwarders by attributes rather than name
- New **Scalability** Constraints.
  - Directory Size.
  - Control Messages.

## Empirical Questions

- How far can Blossom scale...
  - ...before connection latency becomes intolerable for users?
  - ...before directory sizes become unmanageable?
- What can we do to minimize the size and frequency of control messages?
- How can we measure the performance of our transport mechanism (e.g. Tor)?

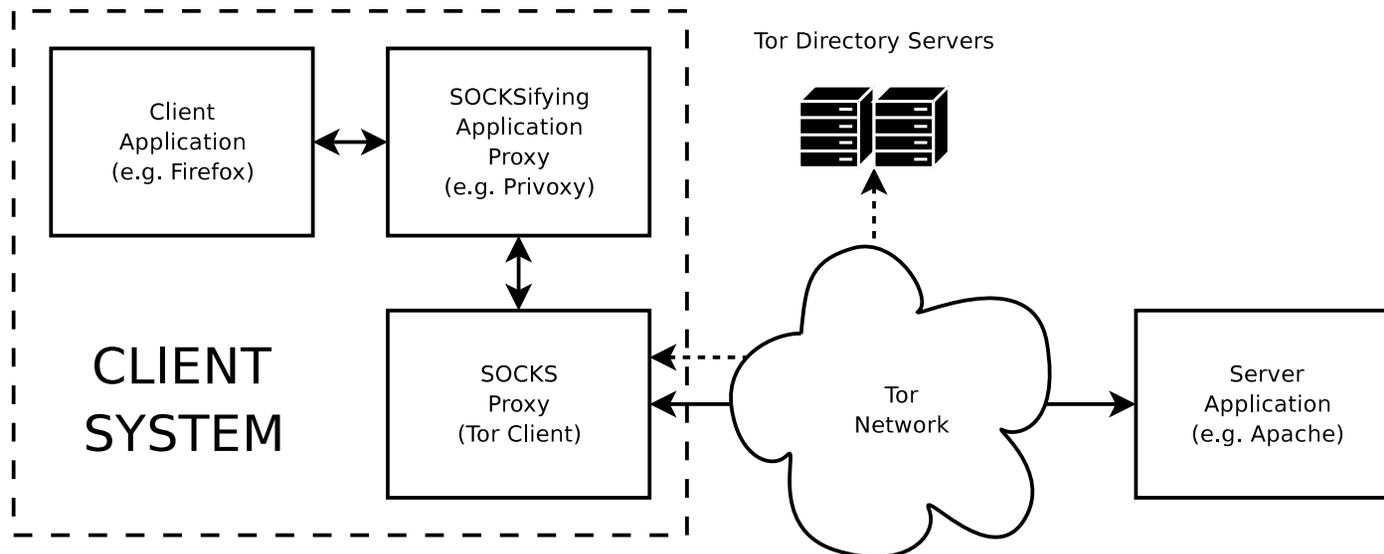
## Philosophical Questions

- What do we lose by building circuits rather than sending independent datagrams?
- How can we use Blossom to enhance Internet search?
- How can Blossom co-exist with reasonable policies established to restrict access to Internet resources?
- How can Blossom co-exist peacefully with unreasonable causes of fragmentation (e.g. oppressive regimes)?

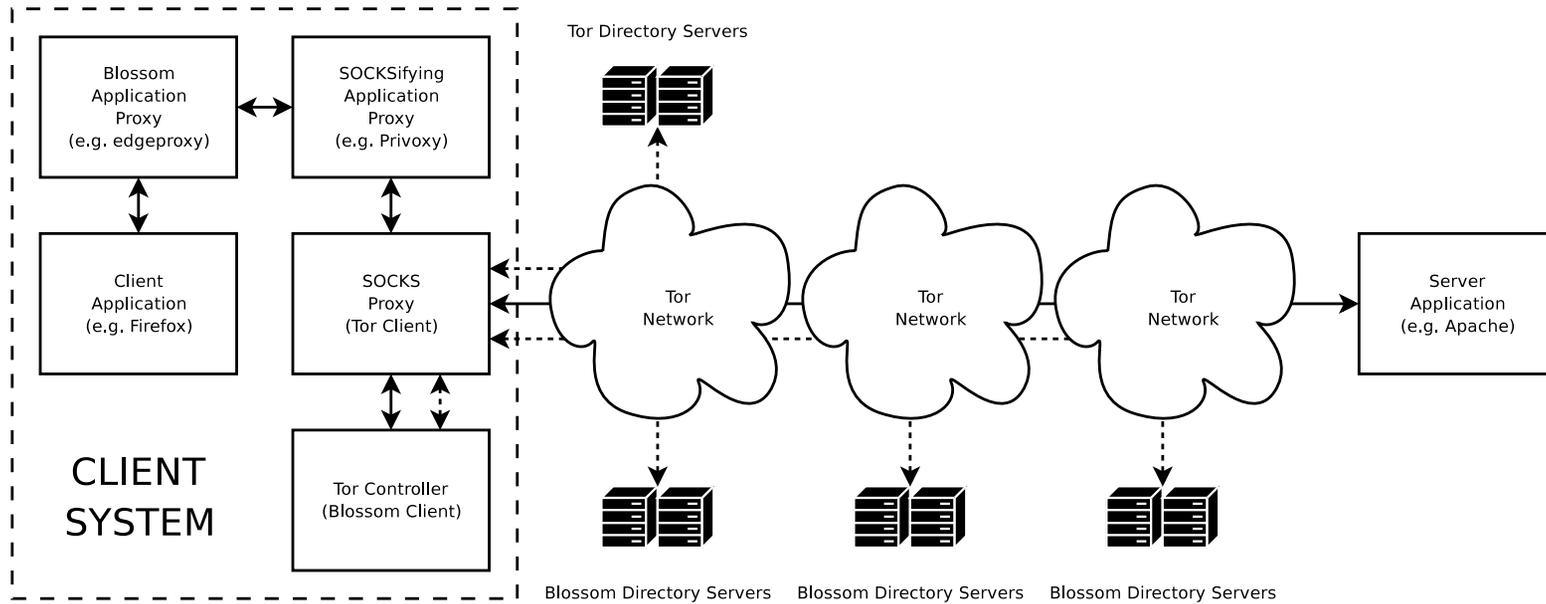


`http://afs.eecs.harvard.edu/~goodell/blossom/`

# Tor Architecture



# Blossom Architecture



# Routing Around Obstacles

- **TRIAD** (Cheriton/Gritter 00)
  - uses globally-unique, hierarchical, DNS-style names to identify networks
  - names propagate through system using BGP-like advertisements
- **RON** (Andersen et al. 01)
  - heal network topology by responding faster than BGP
  - limited to small confederations

## Indirection

- **I3** (Stoica et al. 02)
  - register services with the infrastructure
  - use a commonly-accessible rendezvous point for connectivity
  - substantial inspiration for later work

## Decoupling Policy from Mechanism

- **FARA** (Clark et al. 03)
  - create associations between nodes without requiring global namespace
  - provides a general framework for discovery
- **Platypus** (Snoeren/Raghavan 03)
  - enforce routing policy should on forwarding plane rather than control plane
  - relies upon cooperation from network access providers

## Interoperating with Middleboxes

- **UIP** (Ford 03)
  - argument: hierarchical naming provides efficiency/scalability at the cost of flexibility
  - create single, flat namespace for accessing UIP-enabled resources
- **DOA** (Walfish 04)
  - middleboxes no longer considered harmful
  - create network extension boxes to provide seamless e2e access to DOA resources

## Non-Universal Namespaces

- **Untangling the Web** (Walfish 04)
  - DNS hostnames have special value derived from content provided by web services
  - semantic-free referencing to provide globally-unique self-certifying names
  - client resolves names to semantic-free tags using third-party service, uses resolved tag to access content

## Extending NATs

- **IPNL** (Francis/Gummadi 01)
  - FQDNs as end-system identifiers
  - stateless routers (part of network-layer infrastructure)
  - site-centric: requires configuration and deployment of “frontdoors” between private networks and the core

## Embracing Heterogeneity

- **Plutarch** (Crowcroft et al. 03)
  - coreless
  - acknowledges fragmentation as inevitable
  - requires special configuration of middleboxes to act as gateways between “contexts”
  - resolves names via a peer-to-peer search