

# A System for Authenticated Policy-Compliant Routing

Barath Raghavan and Alex C. Snoeren  
UC San Diego

# Routing Today

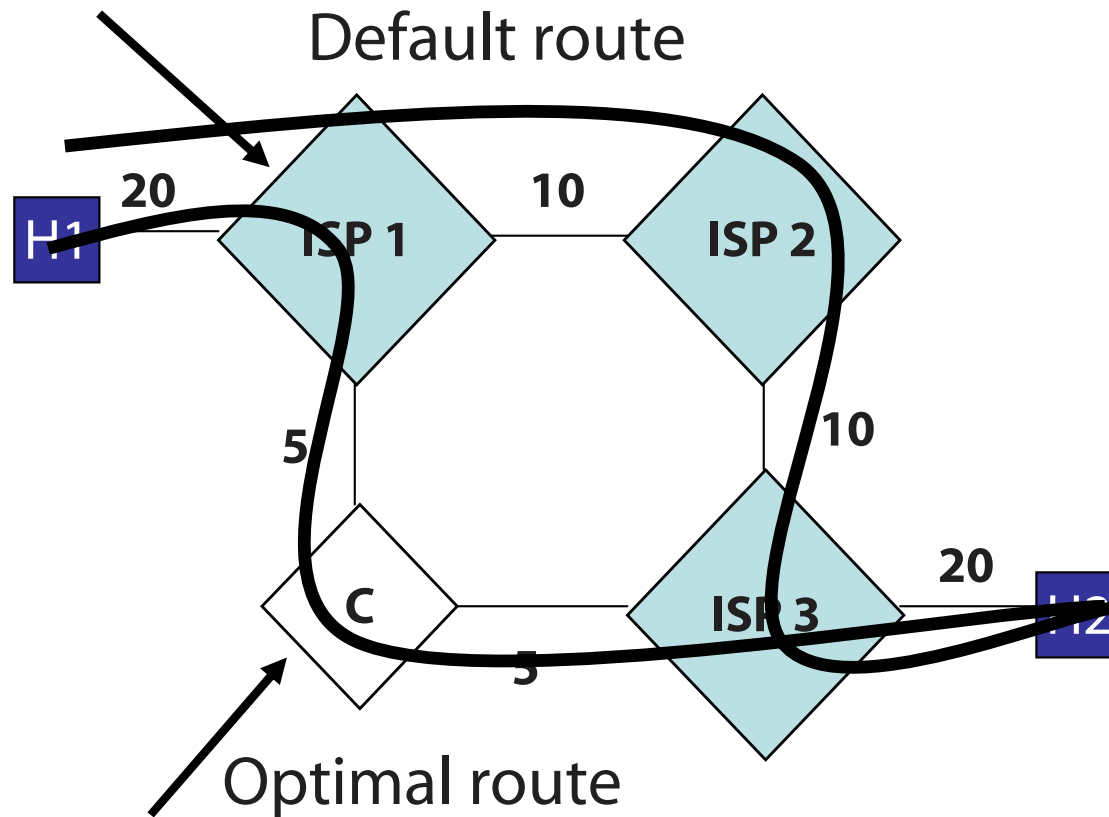
- ISPs perform wide-area routing through BGP
  - Used to express local policy and traffic eng.
  - Problem: users can't express routing preferences
- Overlay routing / IP source routing
  - Enables edge routing control
  - Allows pooling of resources
  - Problem: may interfere with ISP policy and traffic engineering

# Our system: Platypus

- Loose source routing in which...
  - Users can pick their routes
  - ISPs control placement of indirection points
- ... and authentication which enables...
  - ISPs to verify the policy-compliance of traffic
  - Is easily accountable
  - Delegation of source routing rights by users

# An example

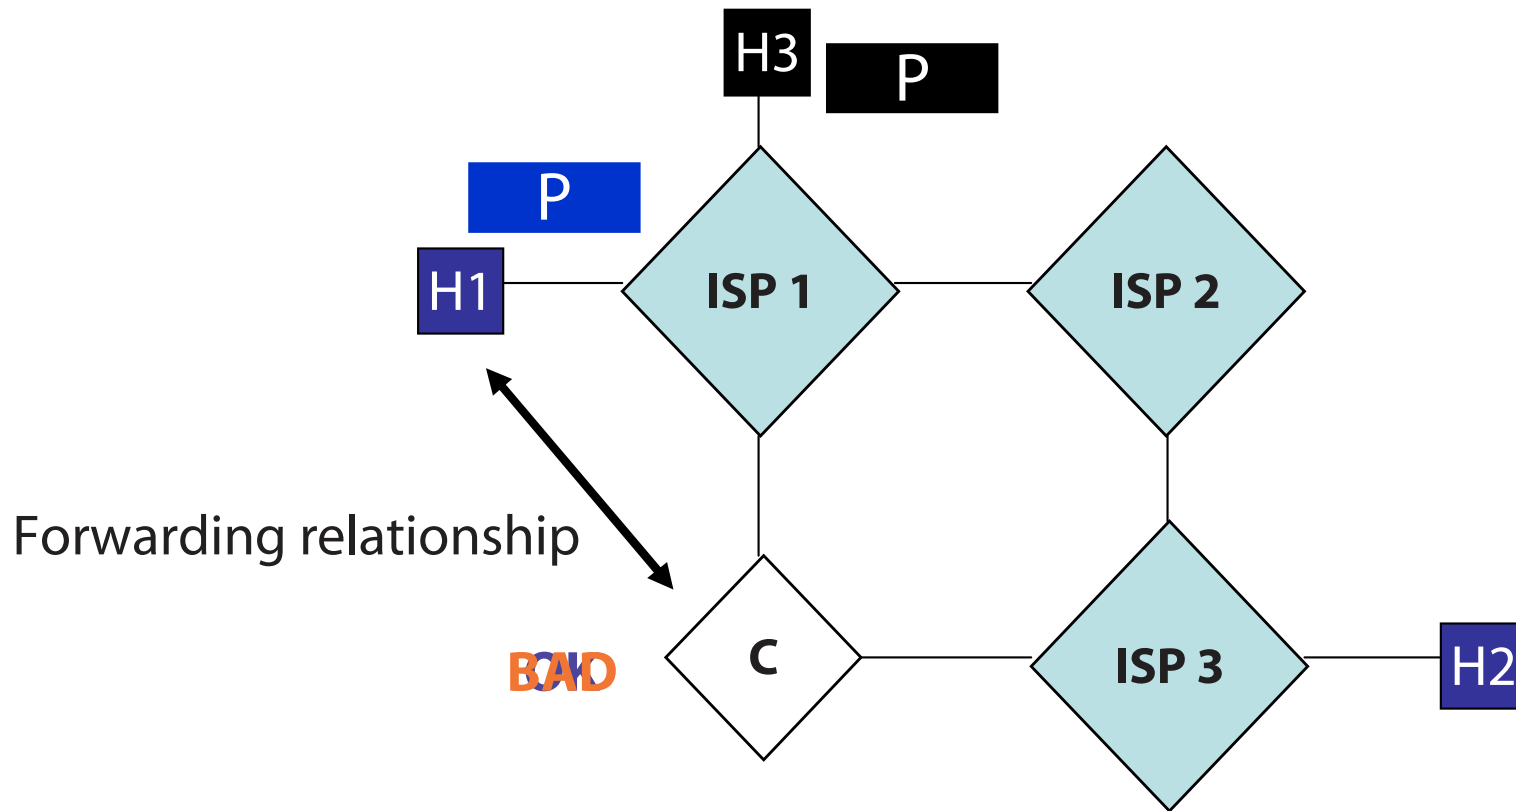
Local policy avoids customer route through **C**



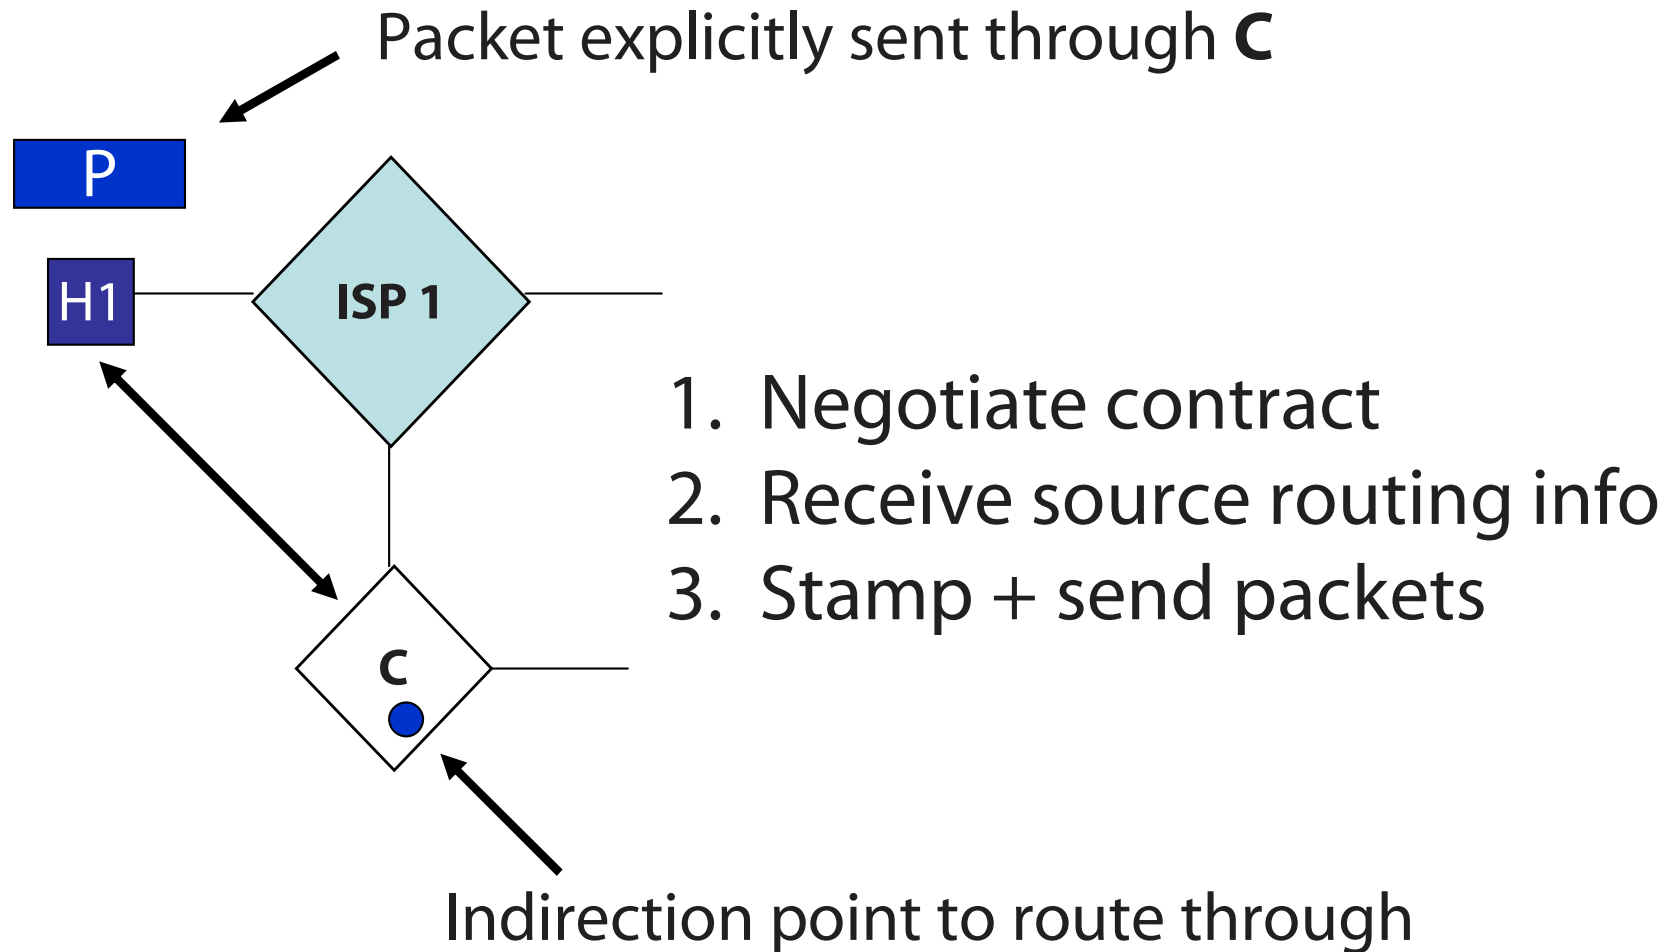
**C** could forward traffic

# The challenge

How can **C** provide forwarding service?



# Our system: Platypus






# Key building blocks

- Routing system providing basic connectivity
- Path discovery mechanisms / services
- Negotiation of business relationships
- Mechanism for authenticated loose source routing

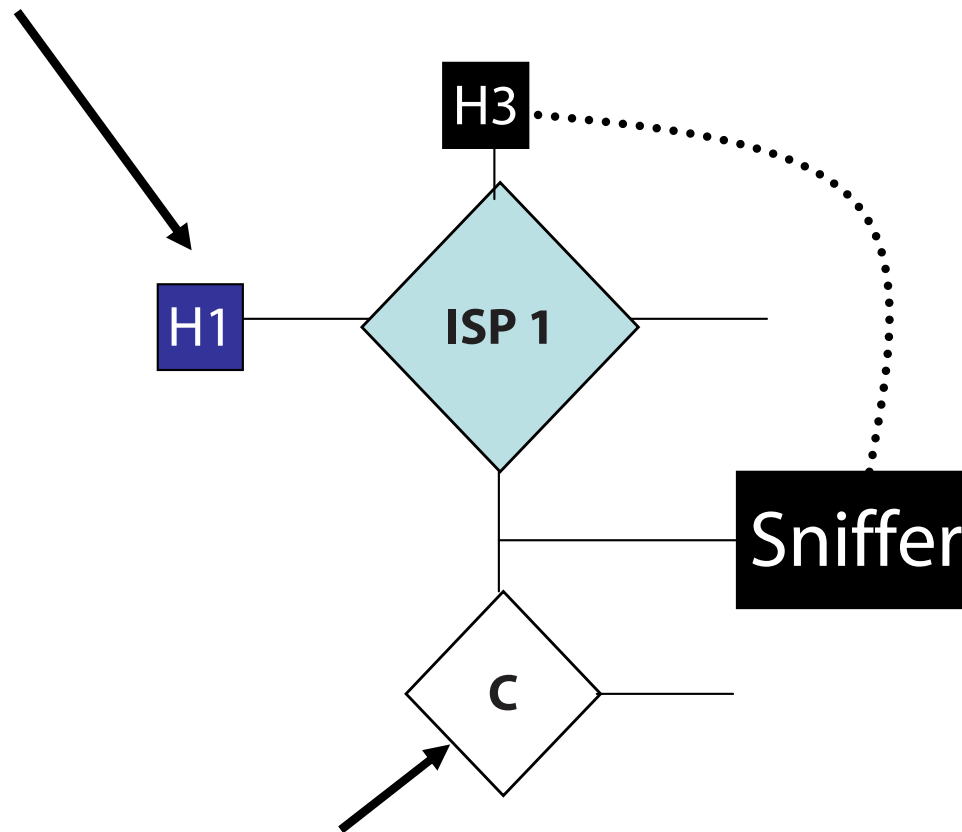
# Network Capabilities

- Specify a hop of the source route, including:
  - Point of indirection, called a *waypoint*
  - The responsible party, called the *resource principal*
- Waypoints are:
  - Chosen by ISPs
  - Specified by a routable IP address



# Authentication

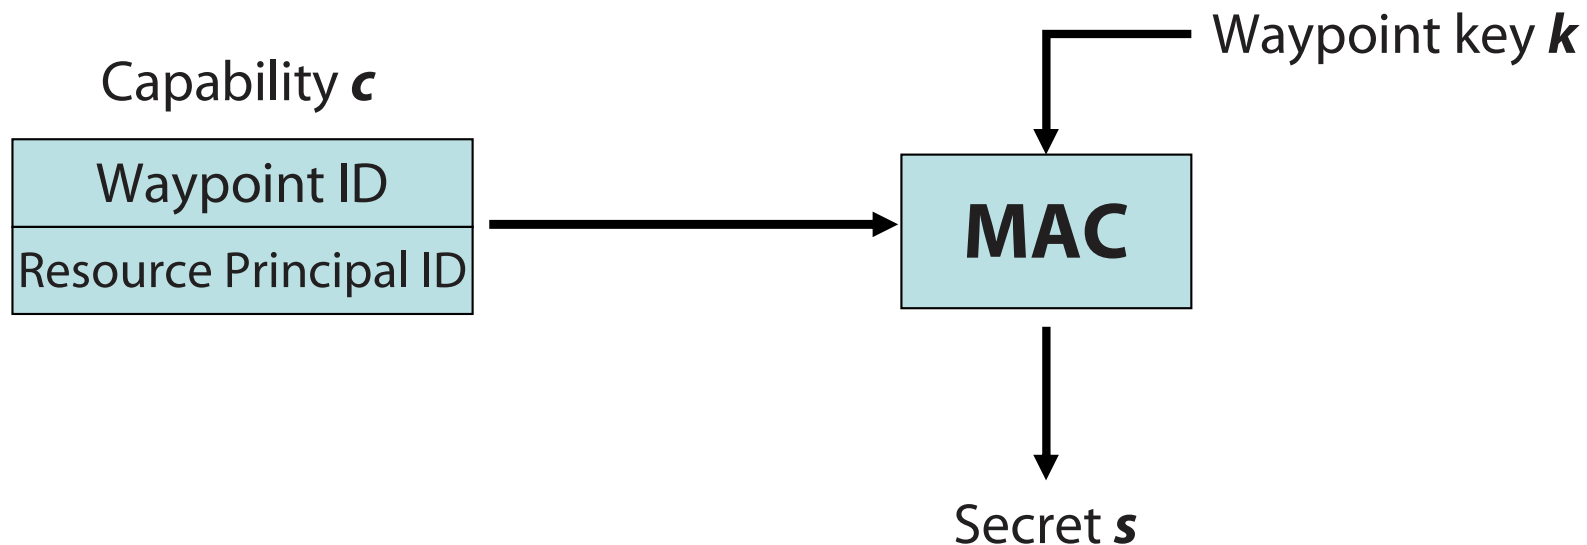
Requires asymmetry of information:  
H1 must know more than H3



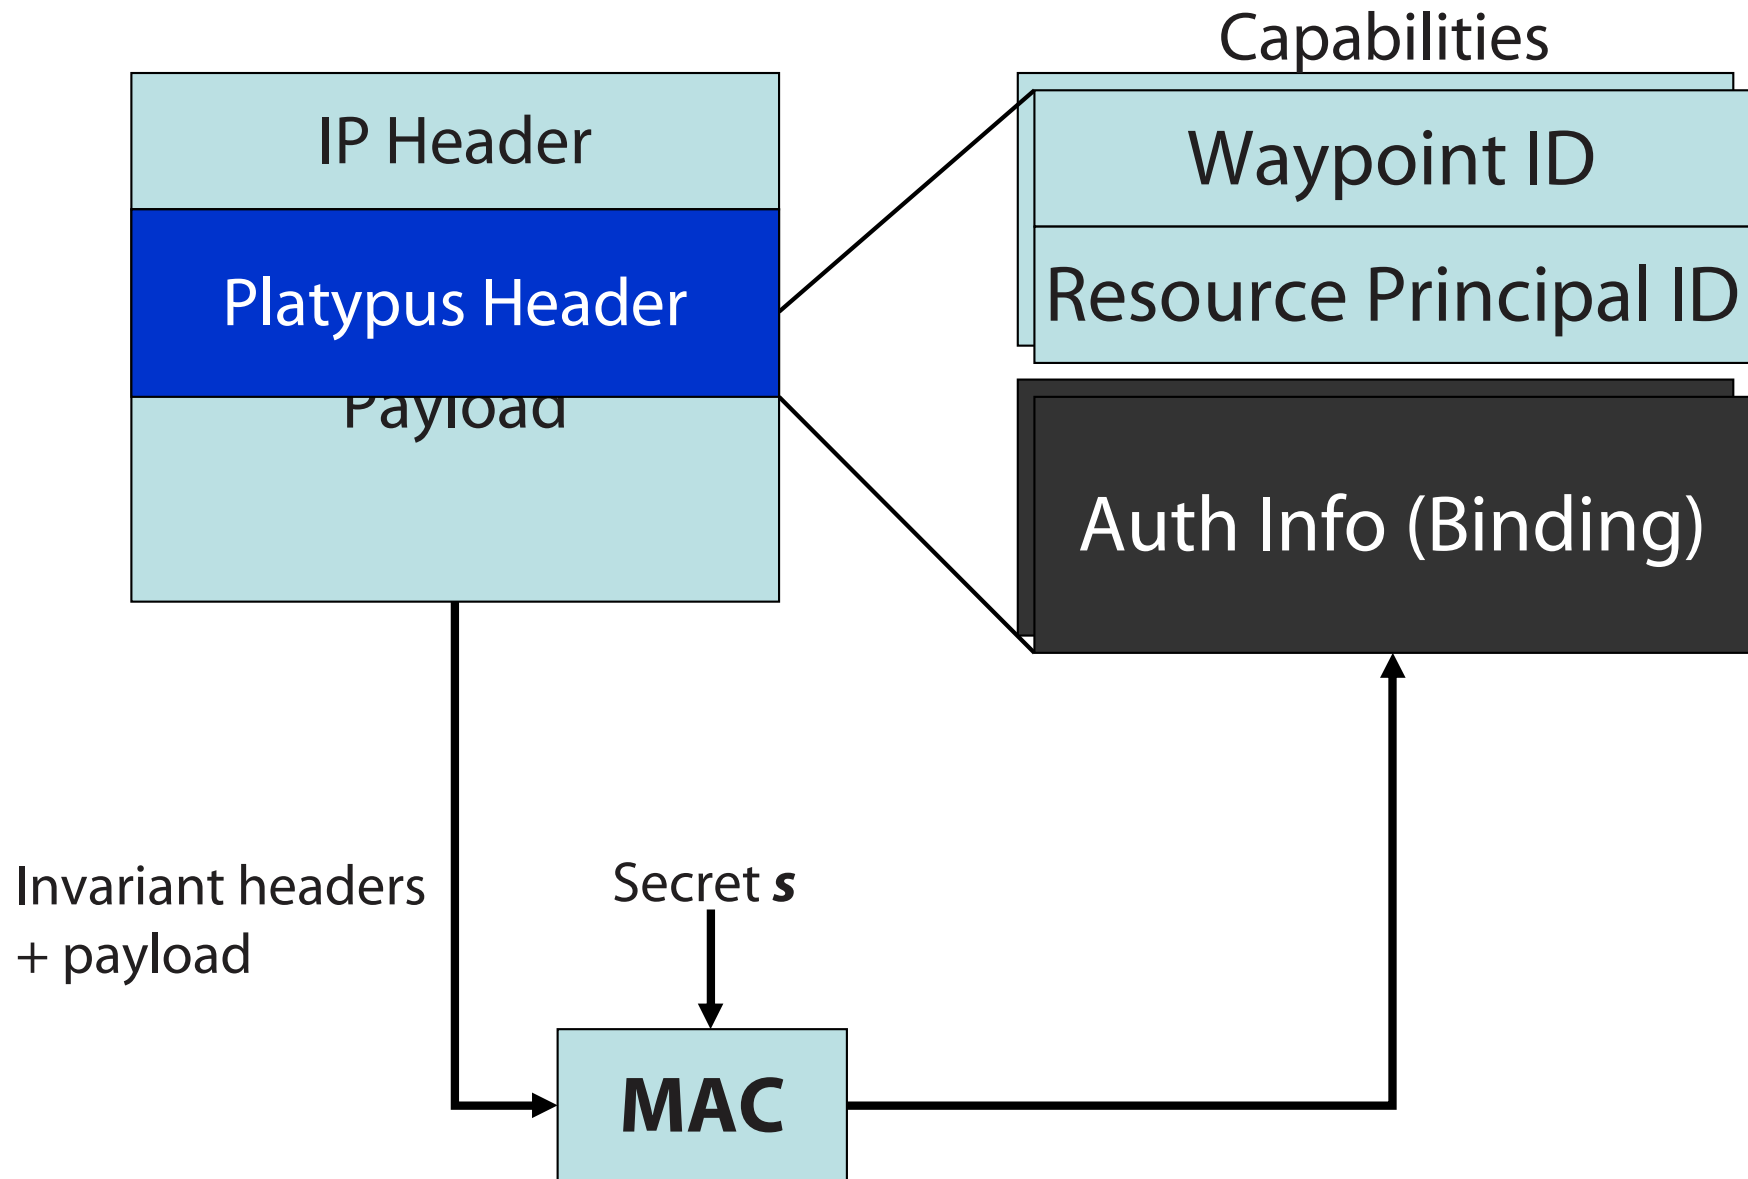
Goal: Distinguish between valid and invalid packets

# Authentication keys

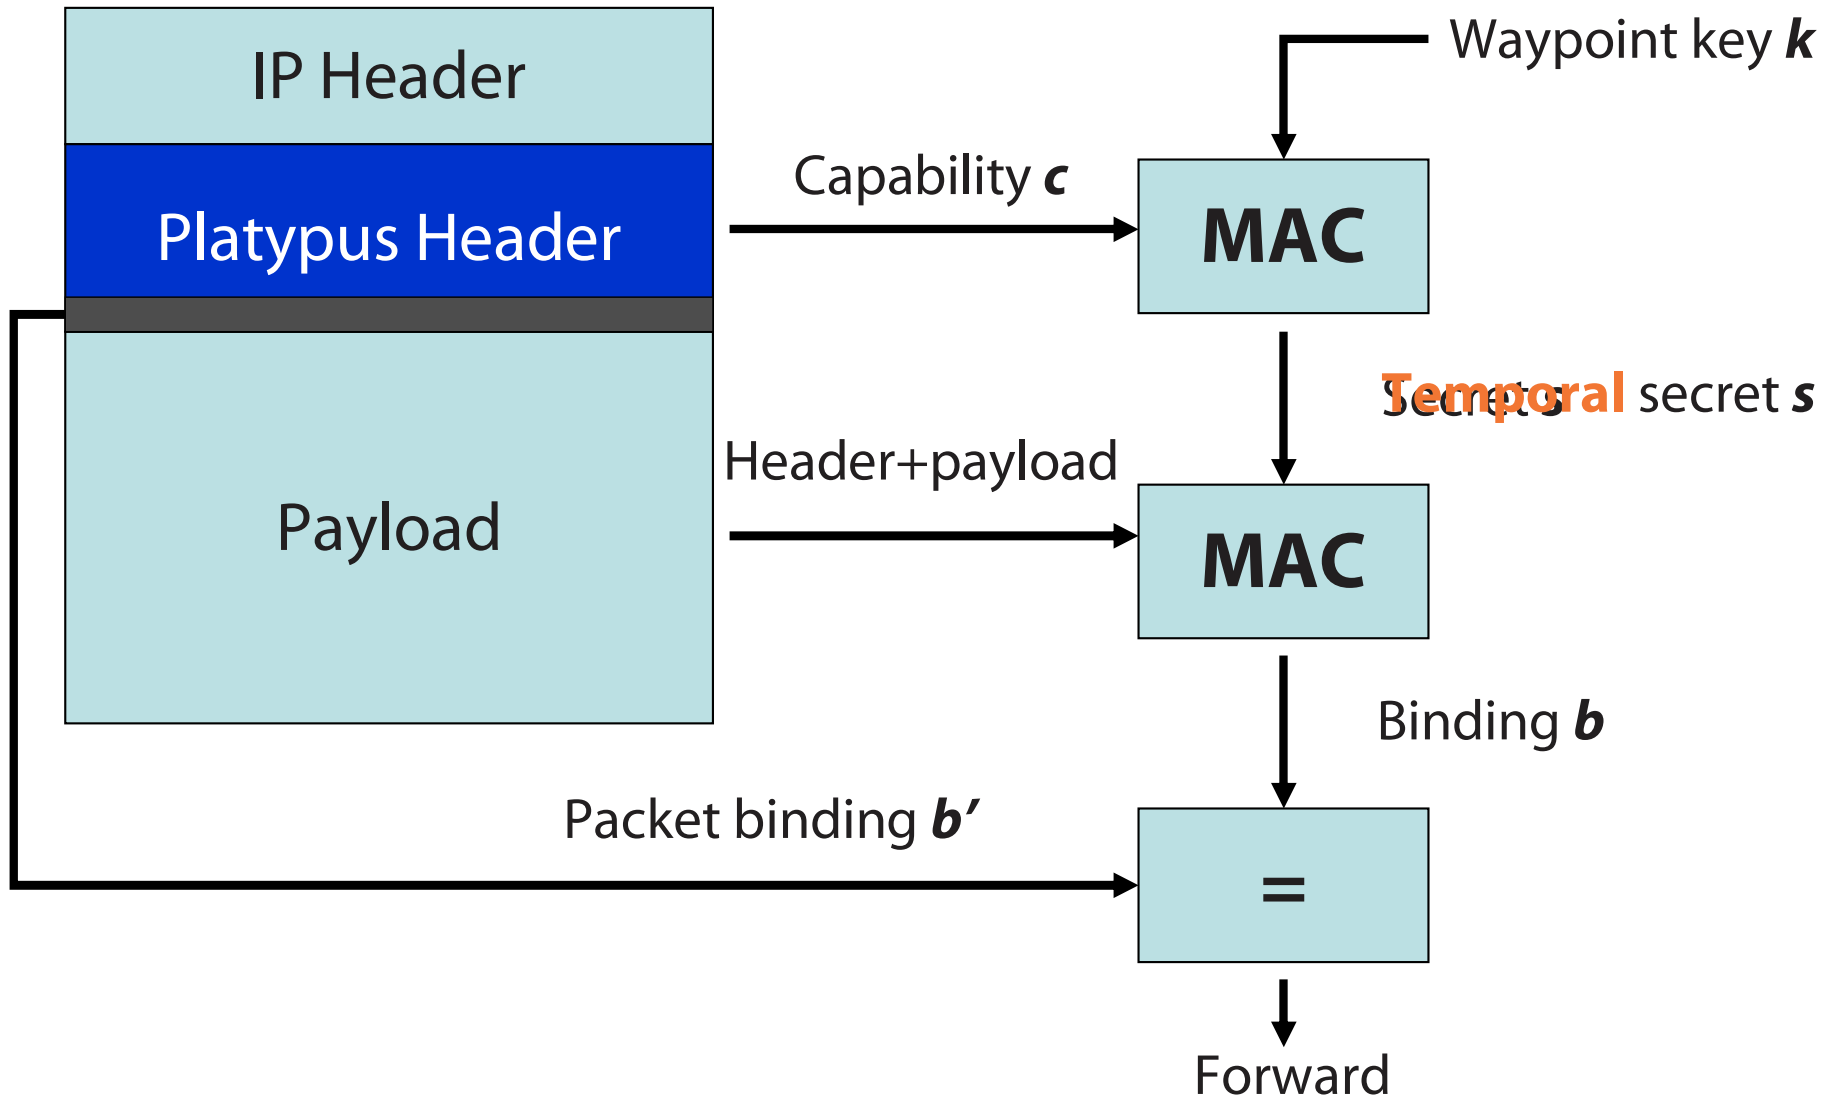
- Each waypoint has one waypoint key  $k$
- Each resource principal has a secret key  $s$ 
  - Derived from waypoint key using a keyed MAC
  - Unique given a waypoint and a capability



# Packet Stamping



# Packet Verification

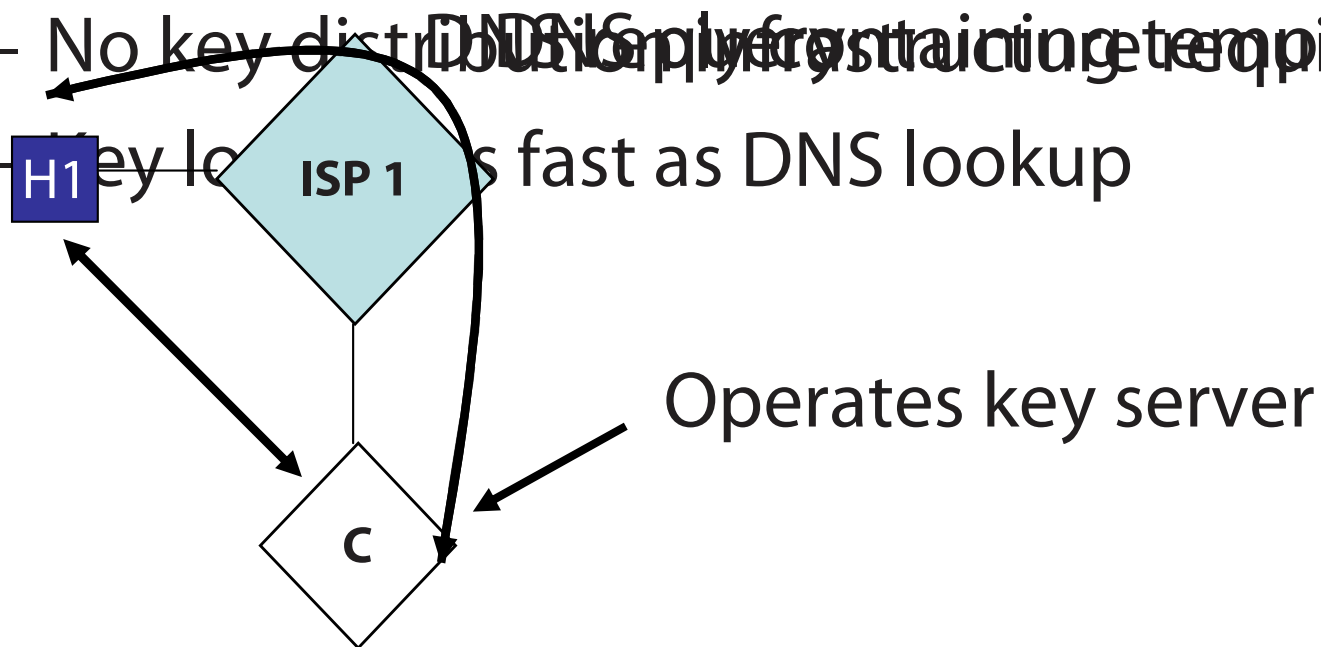


# Temporal secrets

- *Temporal* secret keys expire periodically
  - Expiration allows for changing policies
- No time sync required
  - Secret computation includes Key ID/time
  - Enables expiration on order of clock drift
- Requires lookup of temporal secrets


# Key lookup

- DNS-based key lookup
  - DNS reply contains encrypted secret
  - No key distribution; only for restricting temporal secret
  - Key lookup as fast as DNS lookup



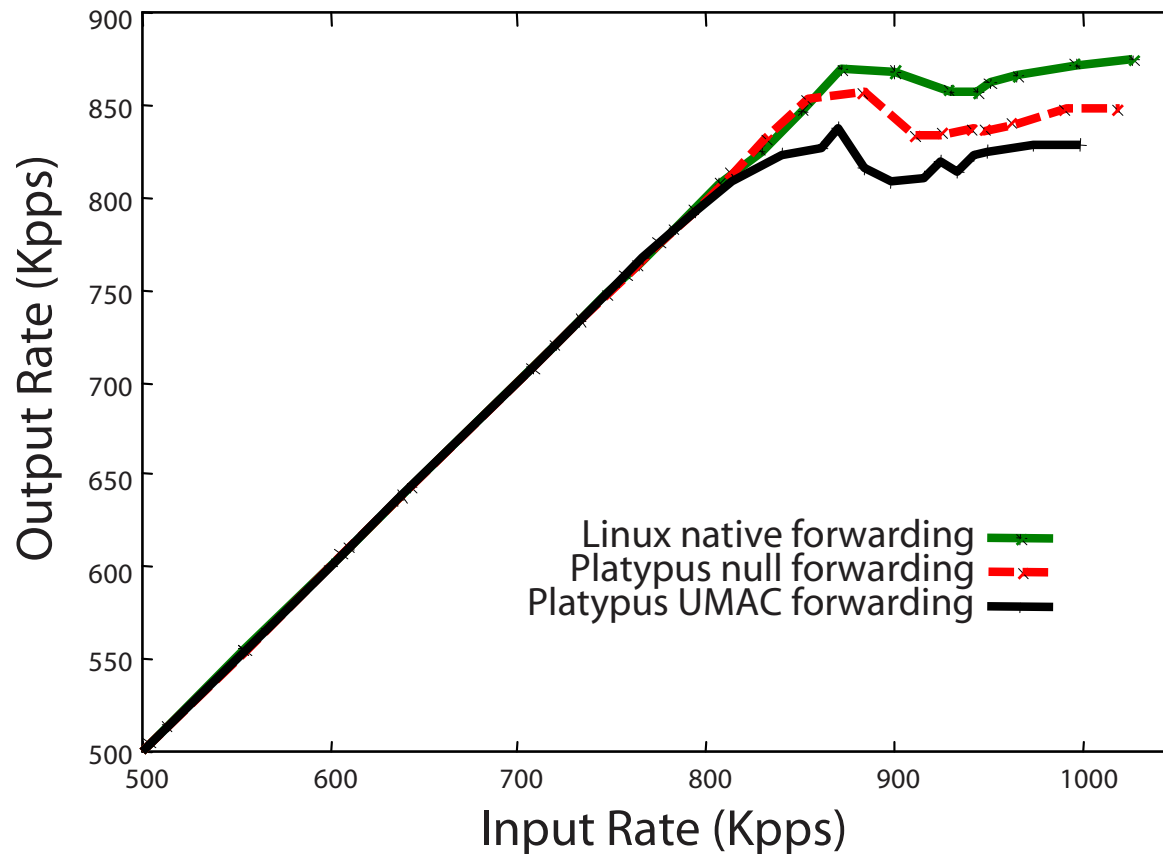
# Delegation

- Users may pass out their capabilities
  - How might they restrict others' use?
- Capability delegation:
  - Principals can restrict capabilities
  - Limits holder to destinations within an IP prefix
  - Useful to ensure similar reverse paths



# Implementation

- End-host based stamping/forwarding
- User-level and kernel module versions



# Per-packet latency

- Total per-packet time = I/O time + header processing
- I/O time  $\sim 2 \mu\text{s}$
- Worst-case header processing time  $< 2 \mu\text{s}$

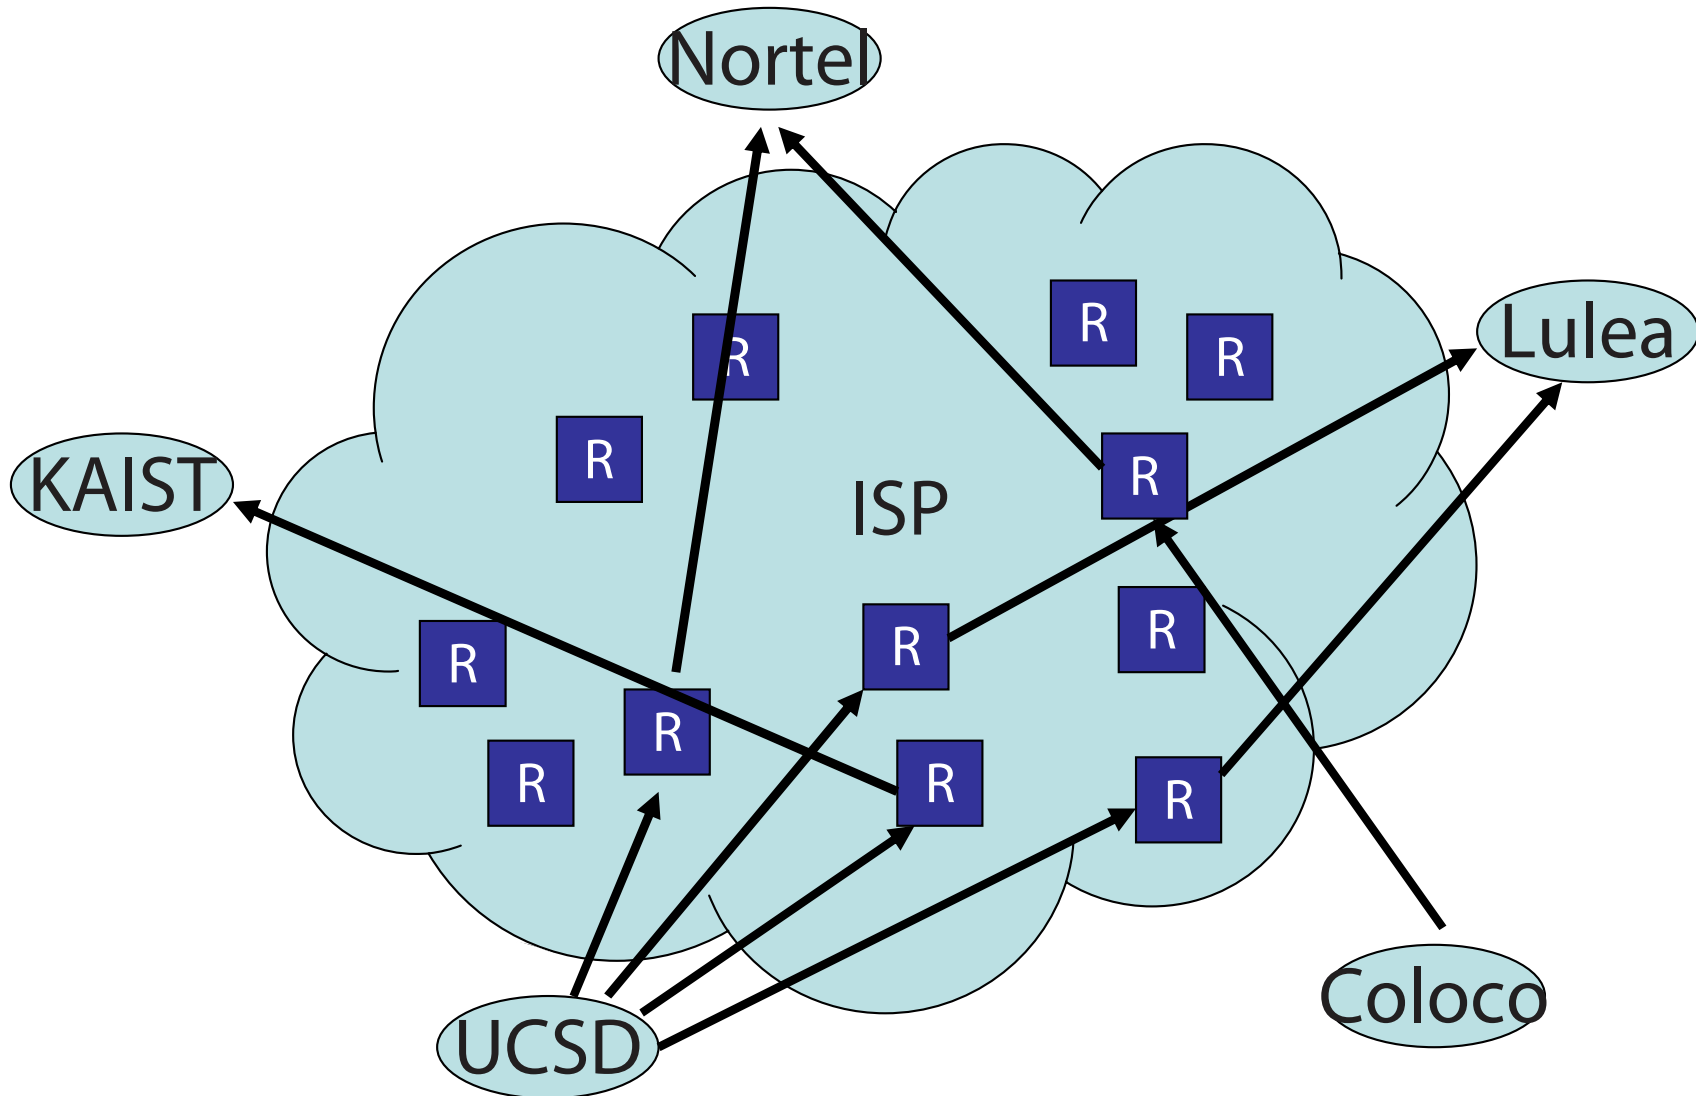
## Header processing overhead

	68 byte	348 byte	1500 byte
Null	172 ns	173 ns	181 ns
UMAC	695 ns	998 ns	1908 ns

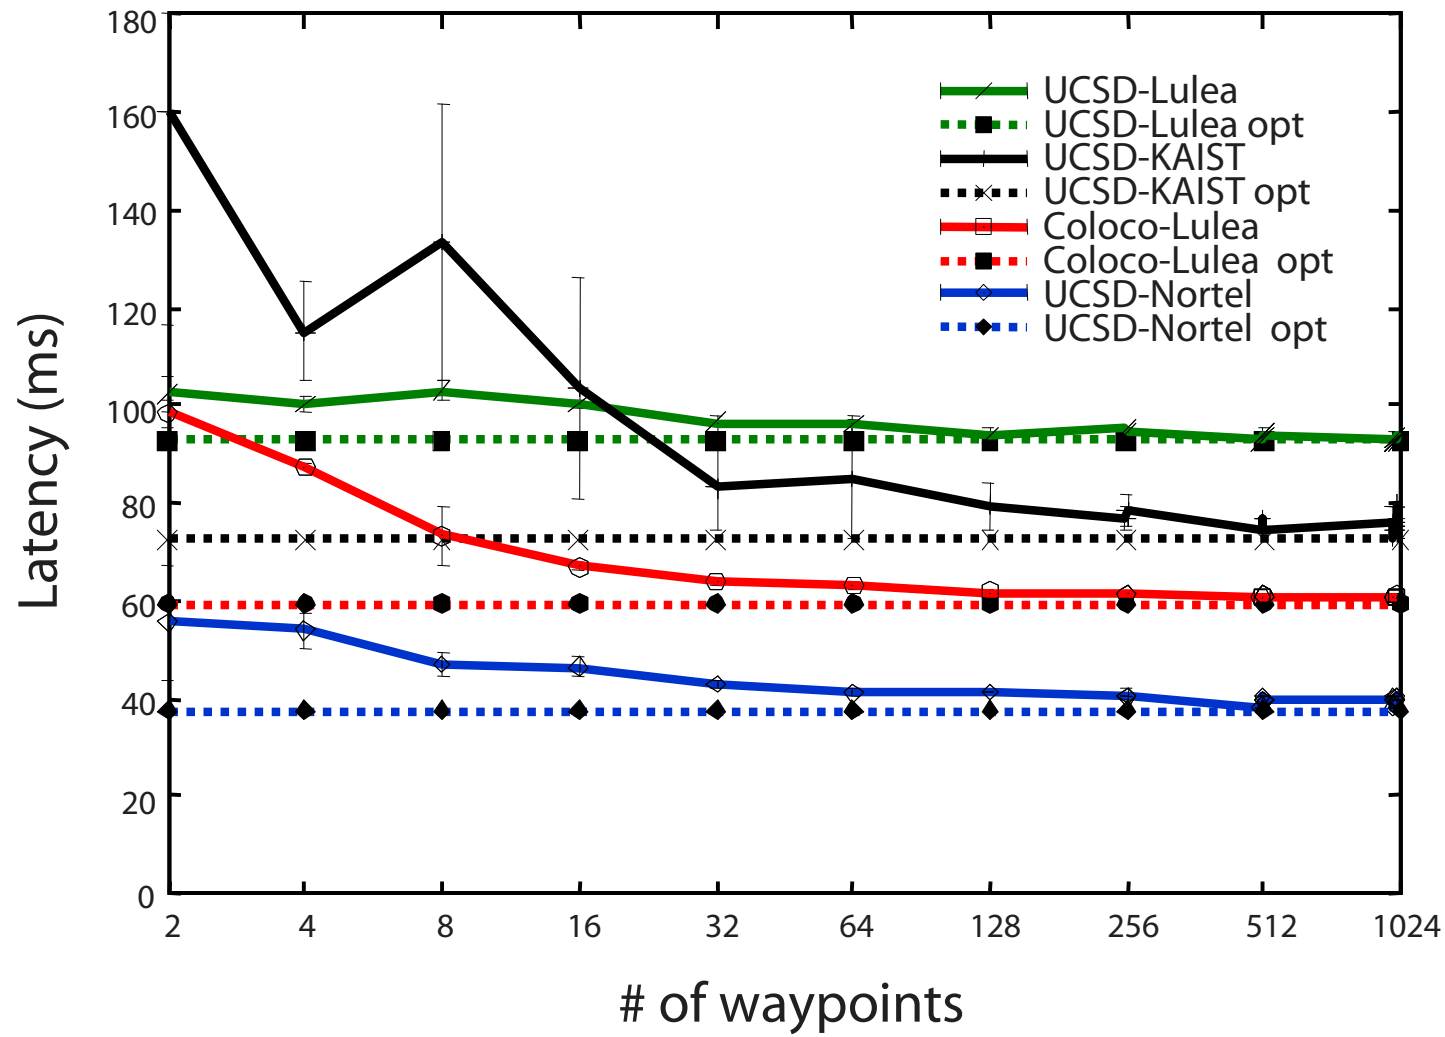
# Deployment

- Incrementally deployable
  - Does not require inter-ISP cooperation
  - *Loose* source-routing based
- How might ISPs deploy Platypus?
  - Where should they be placed?
  - How many Platypus waypoints are needed?

# Measurement study



# Waypoint effectiveness (MCI)



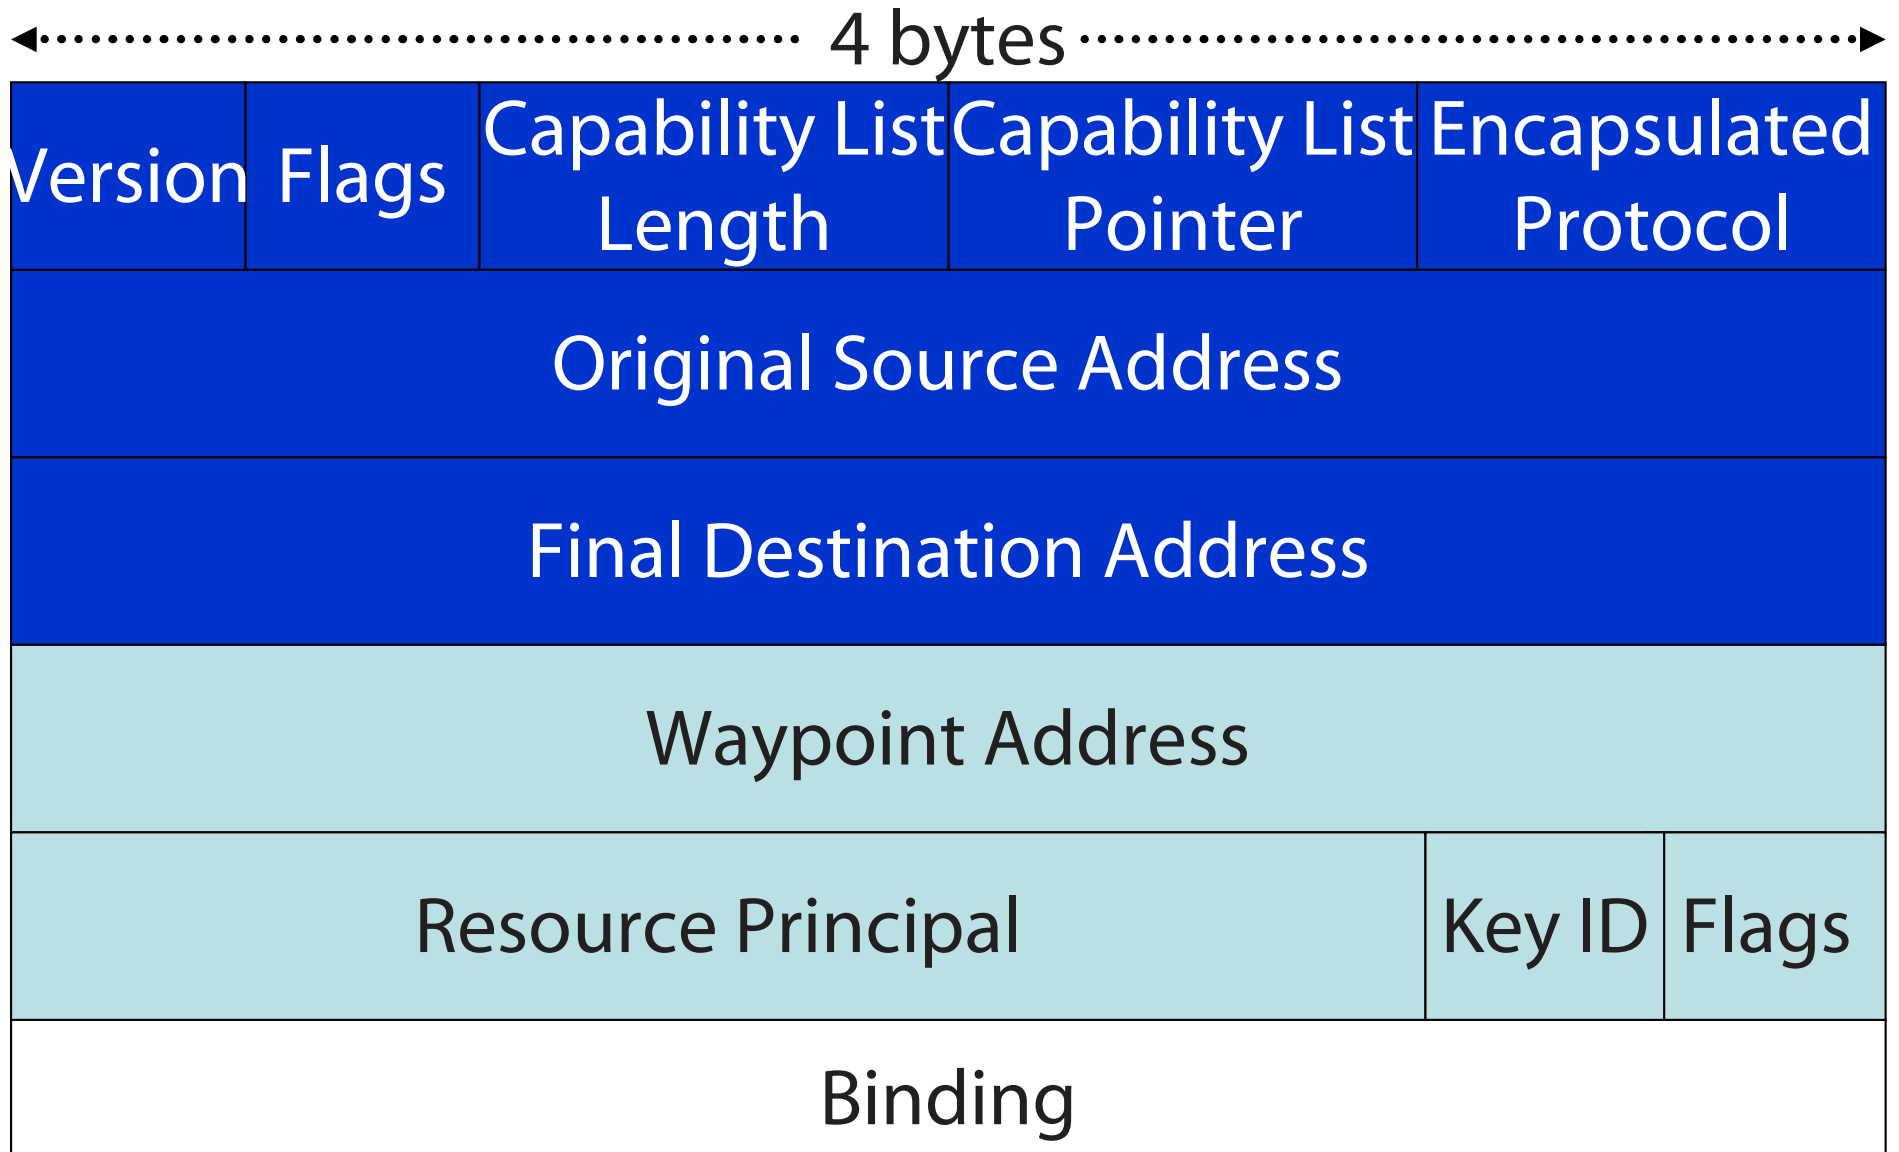
# Summary and future work

- Platypus provides:
  - Source routing with ISP control of waypoints
  - Means for authenticating source routed packets
- Incremental deployment
  - Flow-based Platypus with existing hardware
- New forwarding business model
  - Anyone can sell/resell forwarding service
  - Real-time pricing of capabilities

# Scalability


- Forwarding state
  - Waypoints only need  $O(1)$  state
- Key lookup
  - Lookup overhead is small (3 crypto operations)
  - One key server  $\sim 500,000$  lookups / sec
- Per-principal accounting
  - High speed approx. per-flow counters [Kumar '04]

# Platypus header format




# Temporal secret computation


- For a capability  $c$  and waypoint key  $k$ :  
$$s = \text{MAC}_k(c.\text{way} || c.\text{rp} || (((t \gg n) \& 0\text{xFFFFFFFF0}) | c.\text{id}))$$
- The exception to this is at key ID wraparound
  - $(t \gg n)$  is either incremented or decremented by 1



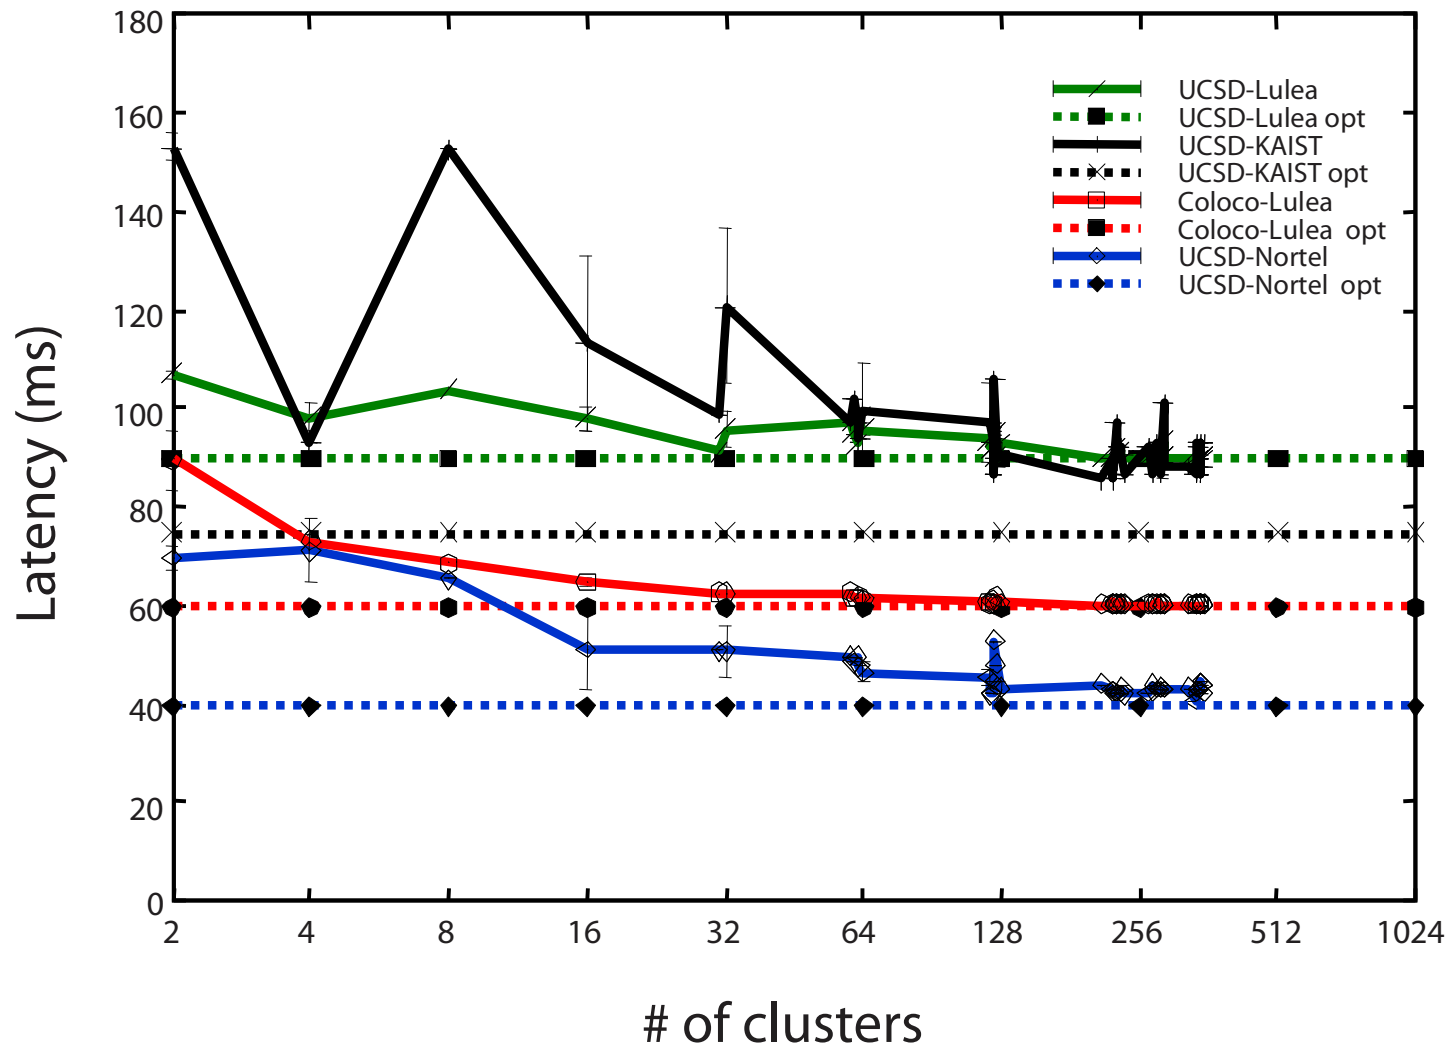
# Measurement results (QWEST)



# Measurement results (GBLX)

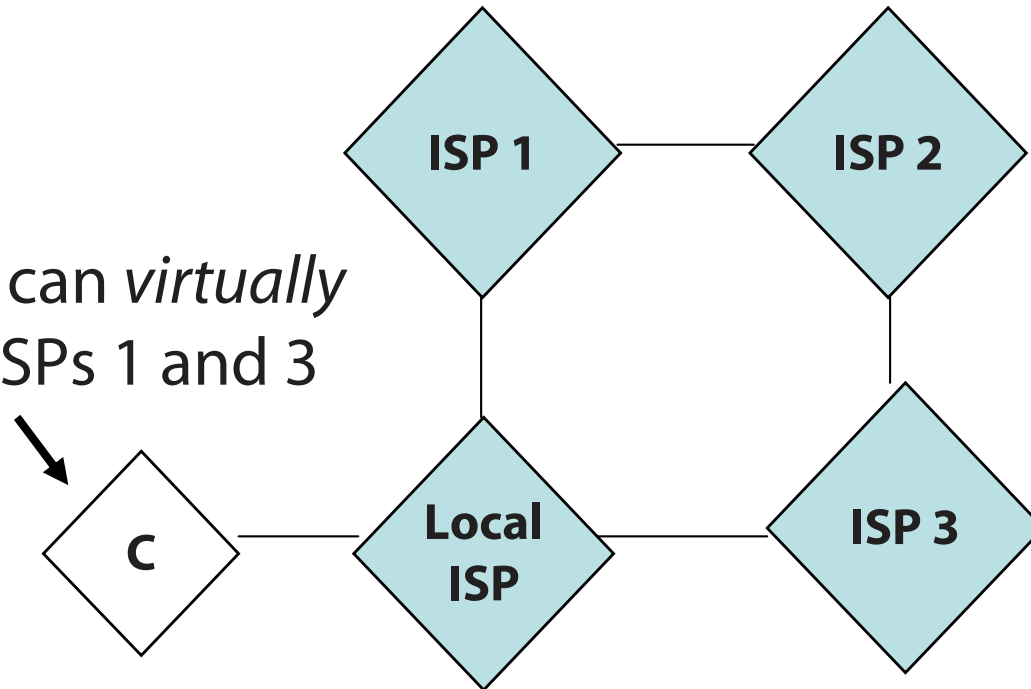


# Measurement results (SPRINT)



# Example: Virtual multihoming

Using Platypus, **C** can *virtually multihome* with ISPs 1 and 3



# Example: Affecting Inbound Traffic

Using Platypus, **C** can distribute delegated capabilities that are restricted to send to prefixes within **C** →

