

# Making Middleboxes Someone Else's Problem: Network Processing as a Cloud Service

Justine Sherry\*, Shaddi Hasan\*,  
Colin Scott\*, Arvind Krishnamurthy†,  
Sylvia Ratnasamy\*, and Vyas Sekar‡



\*




‡




†

# Typical Enterprise Networks




# Typical Enterprise Networks



# A Survey


- 57 enterprise network administrators
- Small (< 1k hosts) to XL (>100k hosts)
- Asked about deployment size, expenses, complexity, and failures.

# How many middleboxes do you deploy?



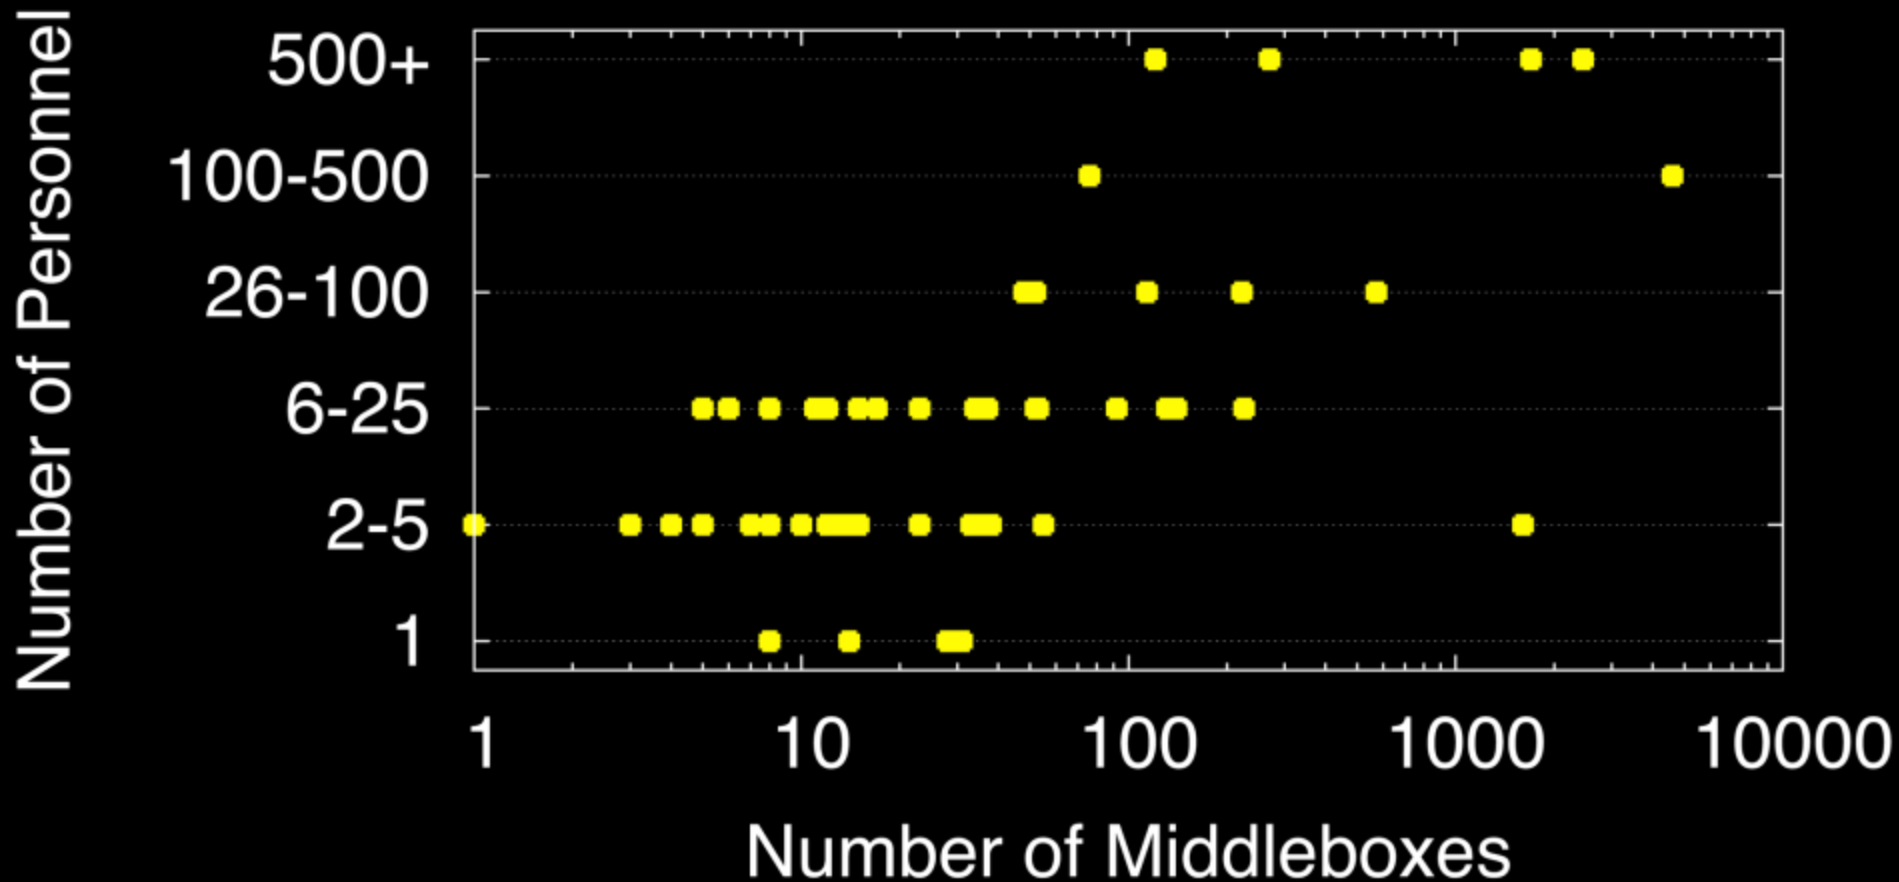
Typically on par with # routers and switches.

# What kinds of middleboxes do you deploy?



Many kinds of devices, all with different functions and management expertise required.

# How many networking personnel are there?



# How do administrators spend their time?

Most administrators spent 1-5 hrs/week dealing with failures; 9% spent 6-10 hrs/week.

	Misconfig.	Overload	Physical/ Electrical
Firewalls	67.3%	16.3%	16.3%
Proxies	63.2%	15.7%	21.1%
IDS	54.45%	11.4%	34%




# Recap


- High Capital and Operating Expenses
- Time Consuming and Error-Prone
- Physical and Overload Failures

How can we improve this?

# Our Proposal



# Our Proposal



# A move to the cloud

- ~~High Capital and Operating Expenses~~
  - Economies of scale and pay-per use
- ~~Time Consuming and Error Prone~~
  - Simplifies configuration and deployment
- ~~Physical and Overload Failures~~
  - Redundant resources for failover

# Our Design

# Challenges

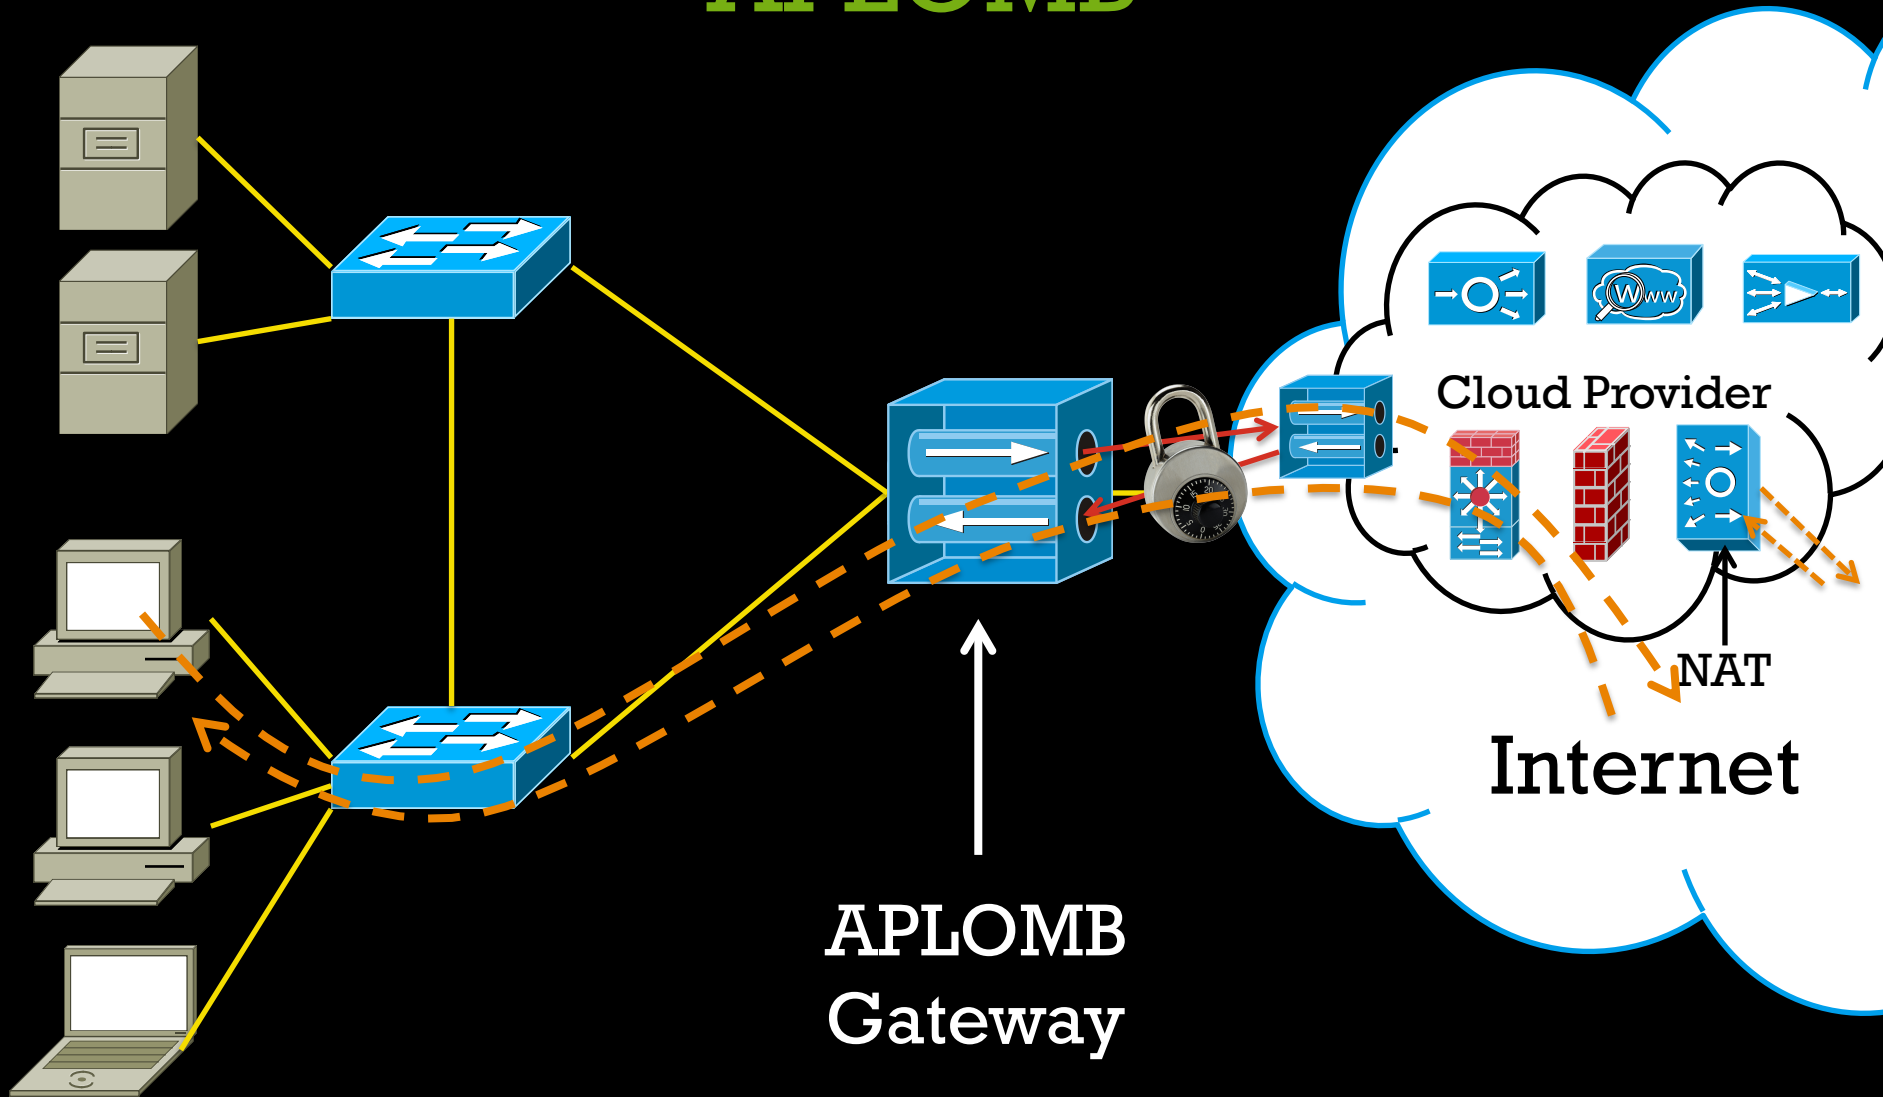
- Minimal Complexity at the Enterprise
- Functional Equivalence
- Low Performance Overhead

**APLOMB**

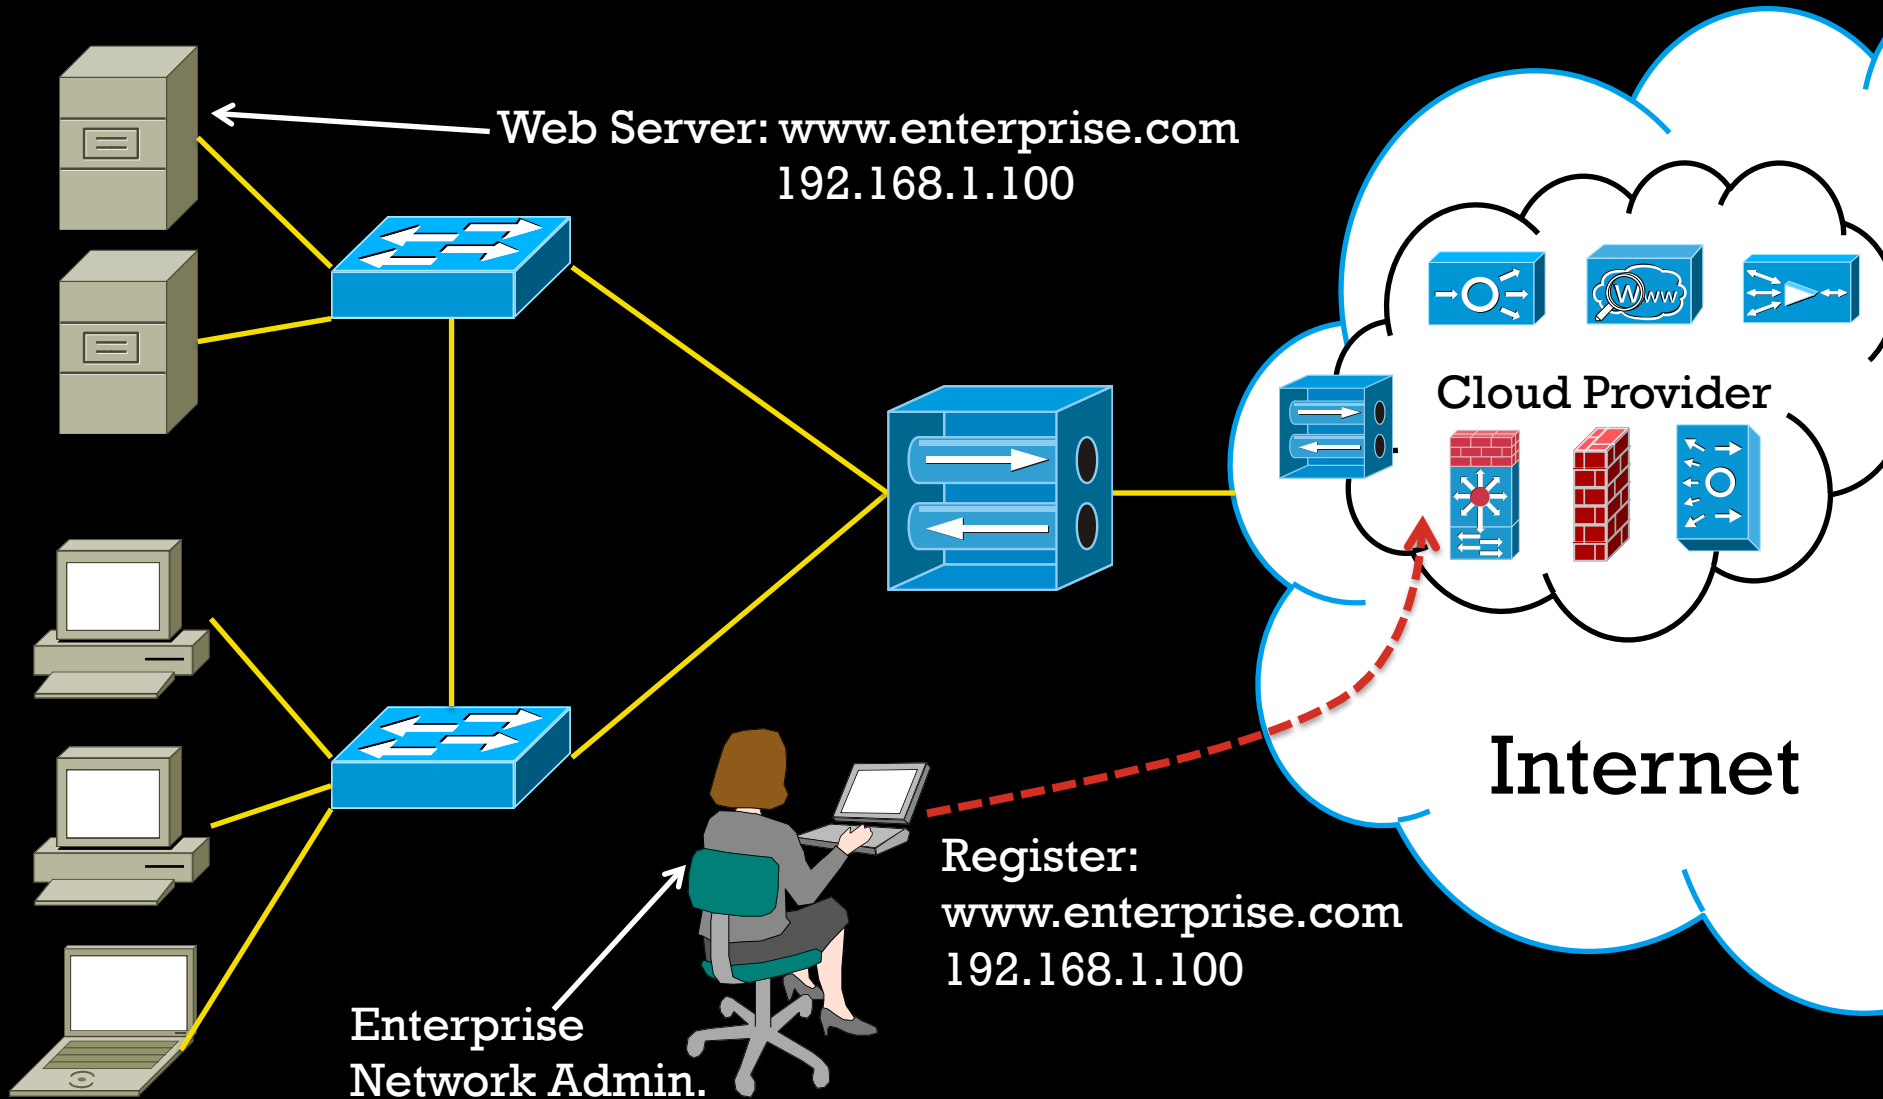
“Appliance for Outsourcing Middleboxes”



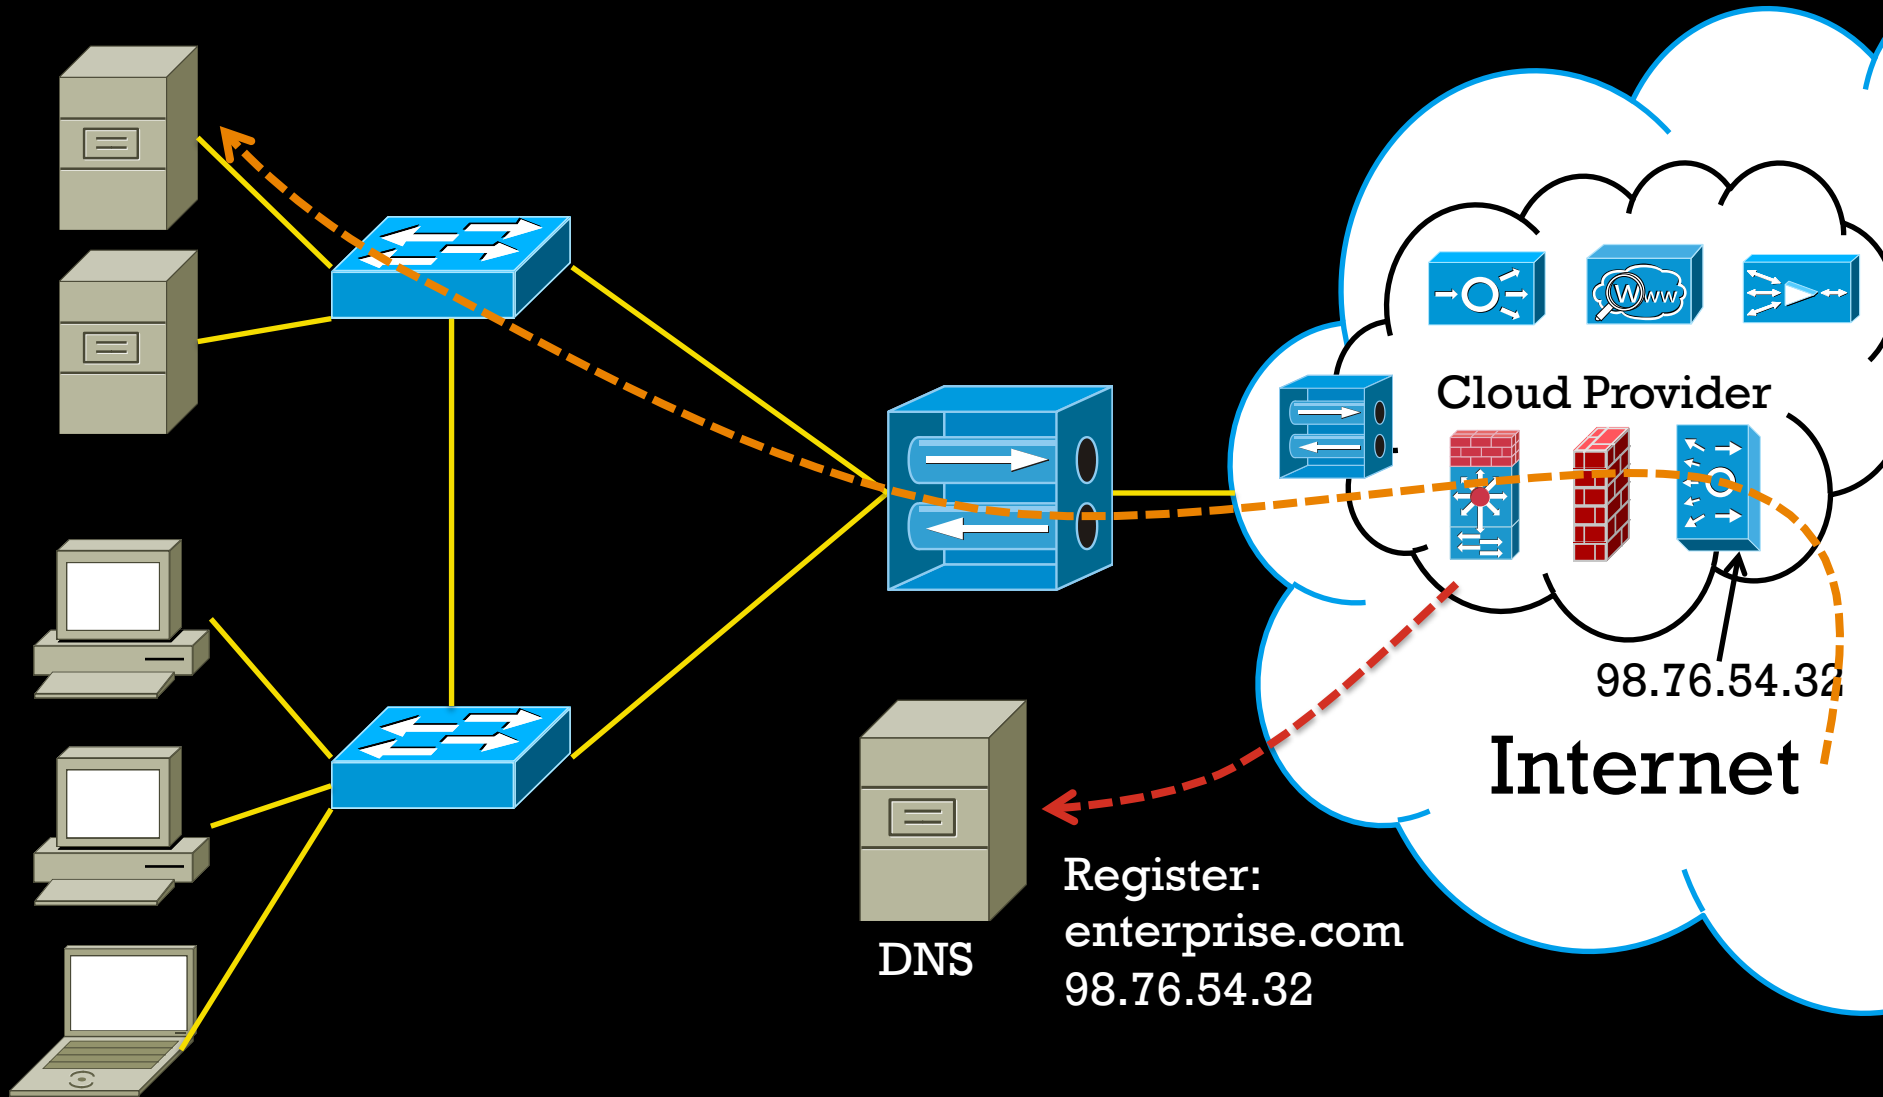
# Outsourcing Middleboxes with APLOMB



# Inbound Traffic

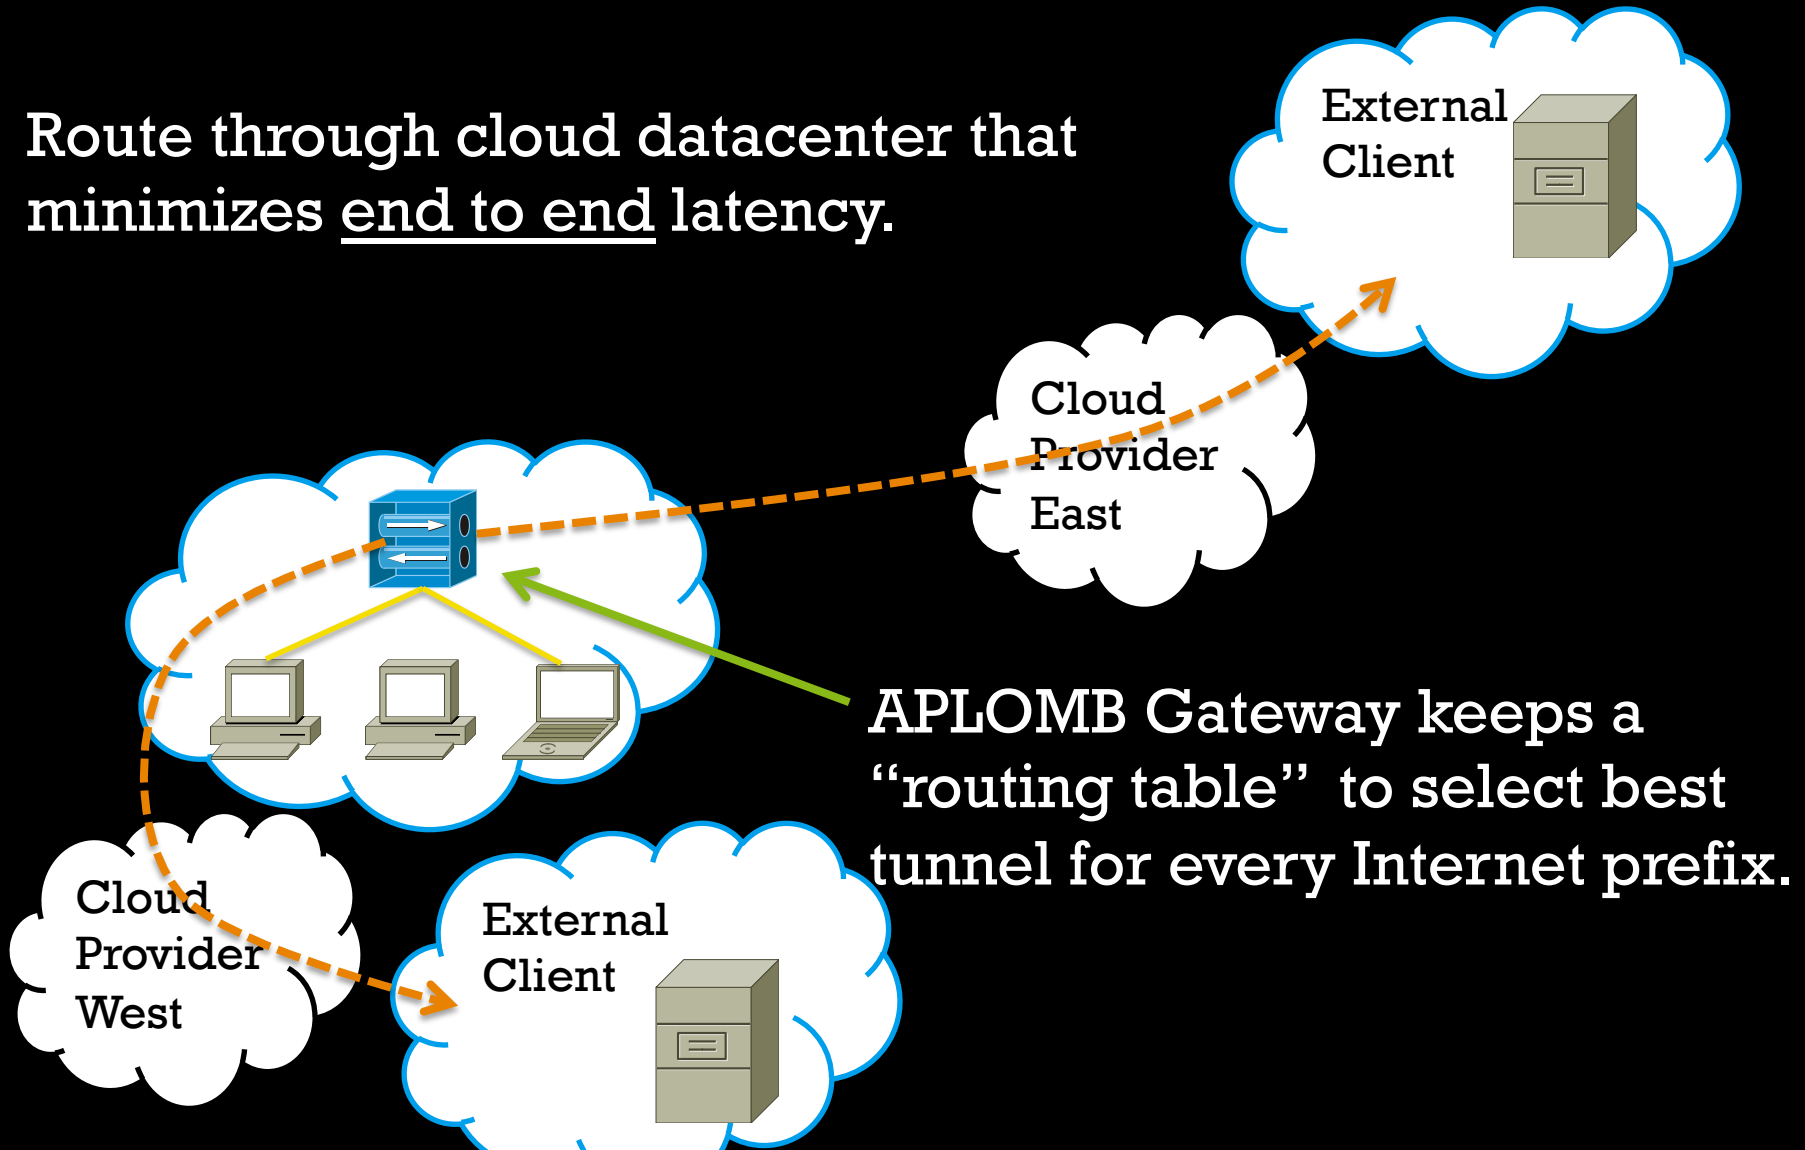


# Inbound Traffic




# Choosing a Datacenter

Route through cloud datacenter that minimizes end to end latency.



# Caches and “Terminal Services”

Traffic destined to services like caches should be redirected to the nearest node.



# APLOMB

“Appliance for Outsourcing Middleboxes”

- Place middleboxes in the cloud.
- Use APLOMB devices and DNS to redirect traffic to and from the cloud.
- That's it.

# Can we outsource all middleboxes?


---

Firewalls	✓
IDSes	✓
Load Balancers	✓
VPNs	✓
Proxy/Caches	✗ Bandwidth?
WAN Optimizers	✗ Compression?

---

# APLOMB+ for Compression

Add generic compression to APLOMB gateway to reduce bandwidth consumption.





# Can we outsource all middleboxes?

---

Firewalls	✓
IDSes	✓
Load Balancers	✓
VPNs	✓
Proxy/Caches	✗ Bandwidth?
WAN Optimizers	✗ Compression?

---



Does it work?

# Our Deployment

- Cloud provider: EC2 – 7 Datacenters
- OpenVPN for tunneling, Vyatta for middlebox services
- Two Types of Clients:
  - Software VPN client on laptops
  - Tunneling software router for wired hosts

# Three Part Evaluation

## Implementation & Deployment

- Performance metrics

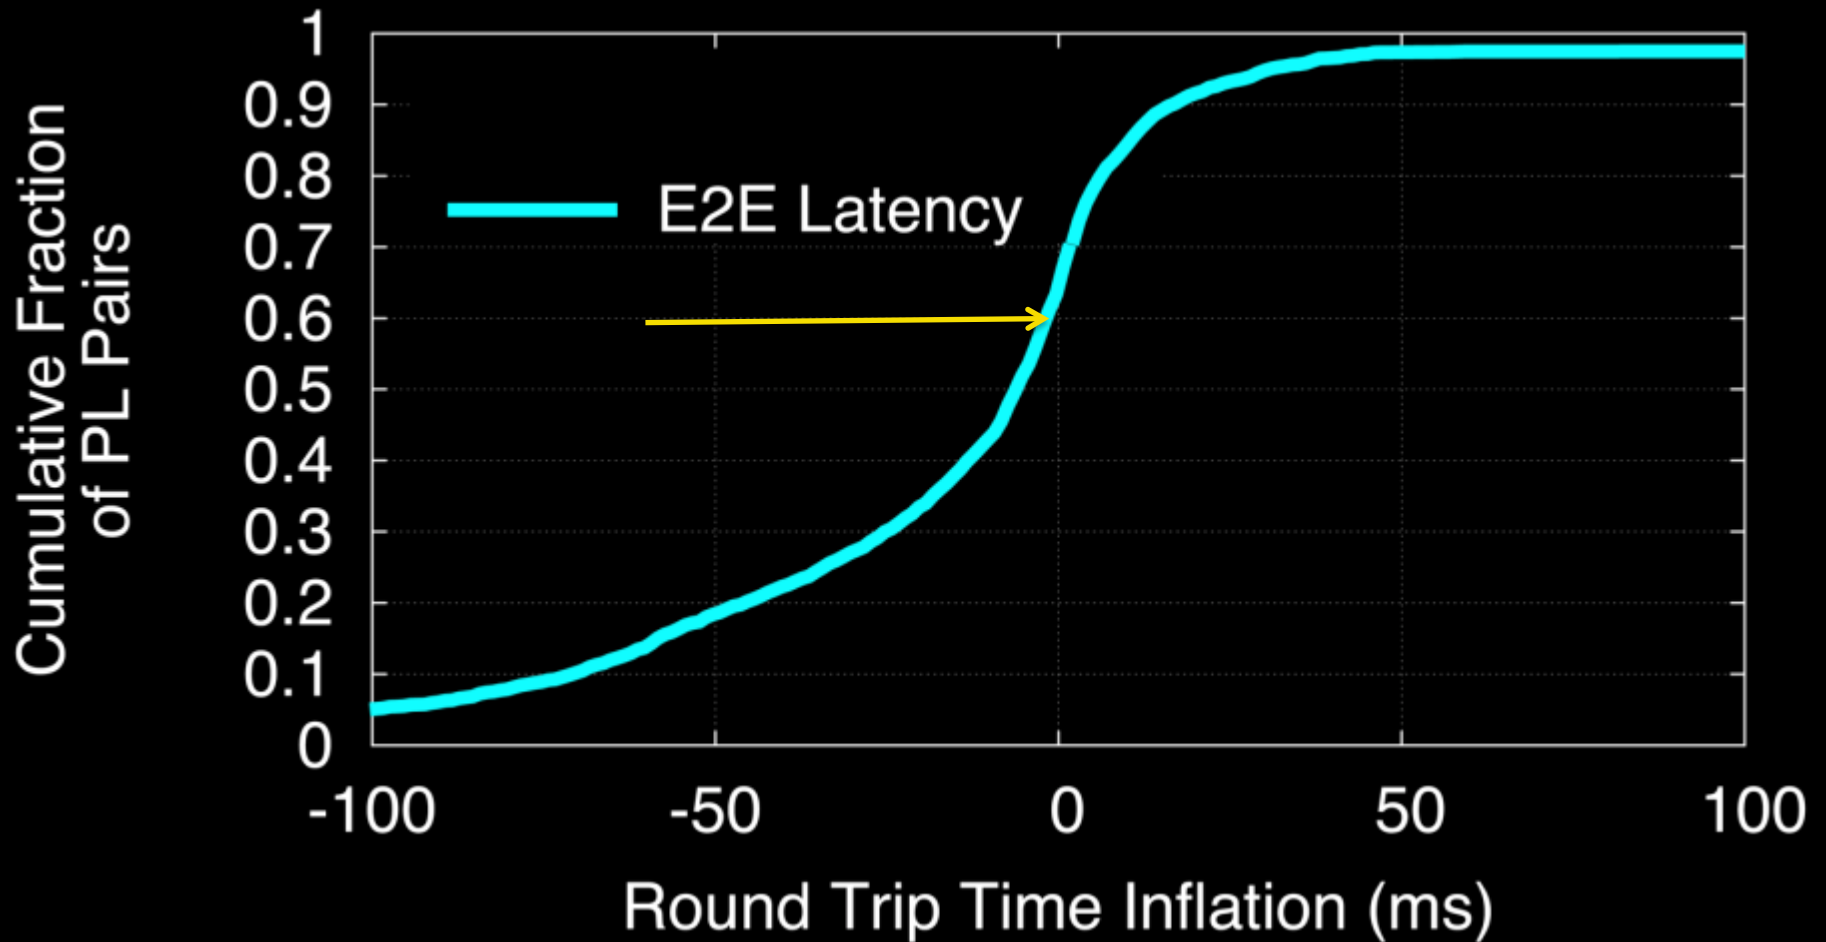
## Wide-Area Measurements

- Network latency

## Case Study of a Large Enterprise

- Impact in a real usage scenario

Does APLOMB inflate latency?



For PlanetLab nodes, 60% of pairs' latency improves with redirection through EC2.

# Latency at a Large Enterprise

Measured redirection latency between enterprise sites.


- Median latency inflation: 1.13 ms
- Sites experiencing inflation were primarily in areas where EC2 does not have a wide footprint.

How does APLOMB impact other quality metrics, like bandwidth and jitter?



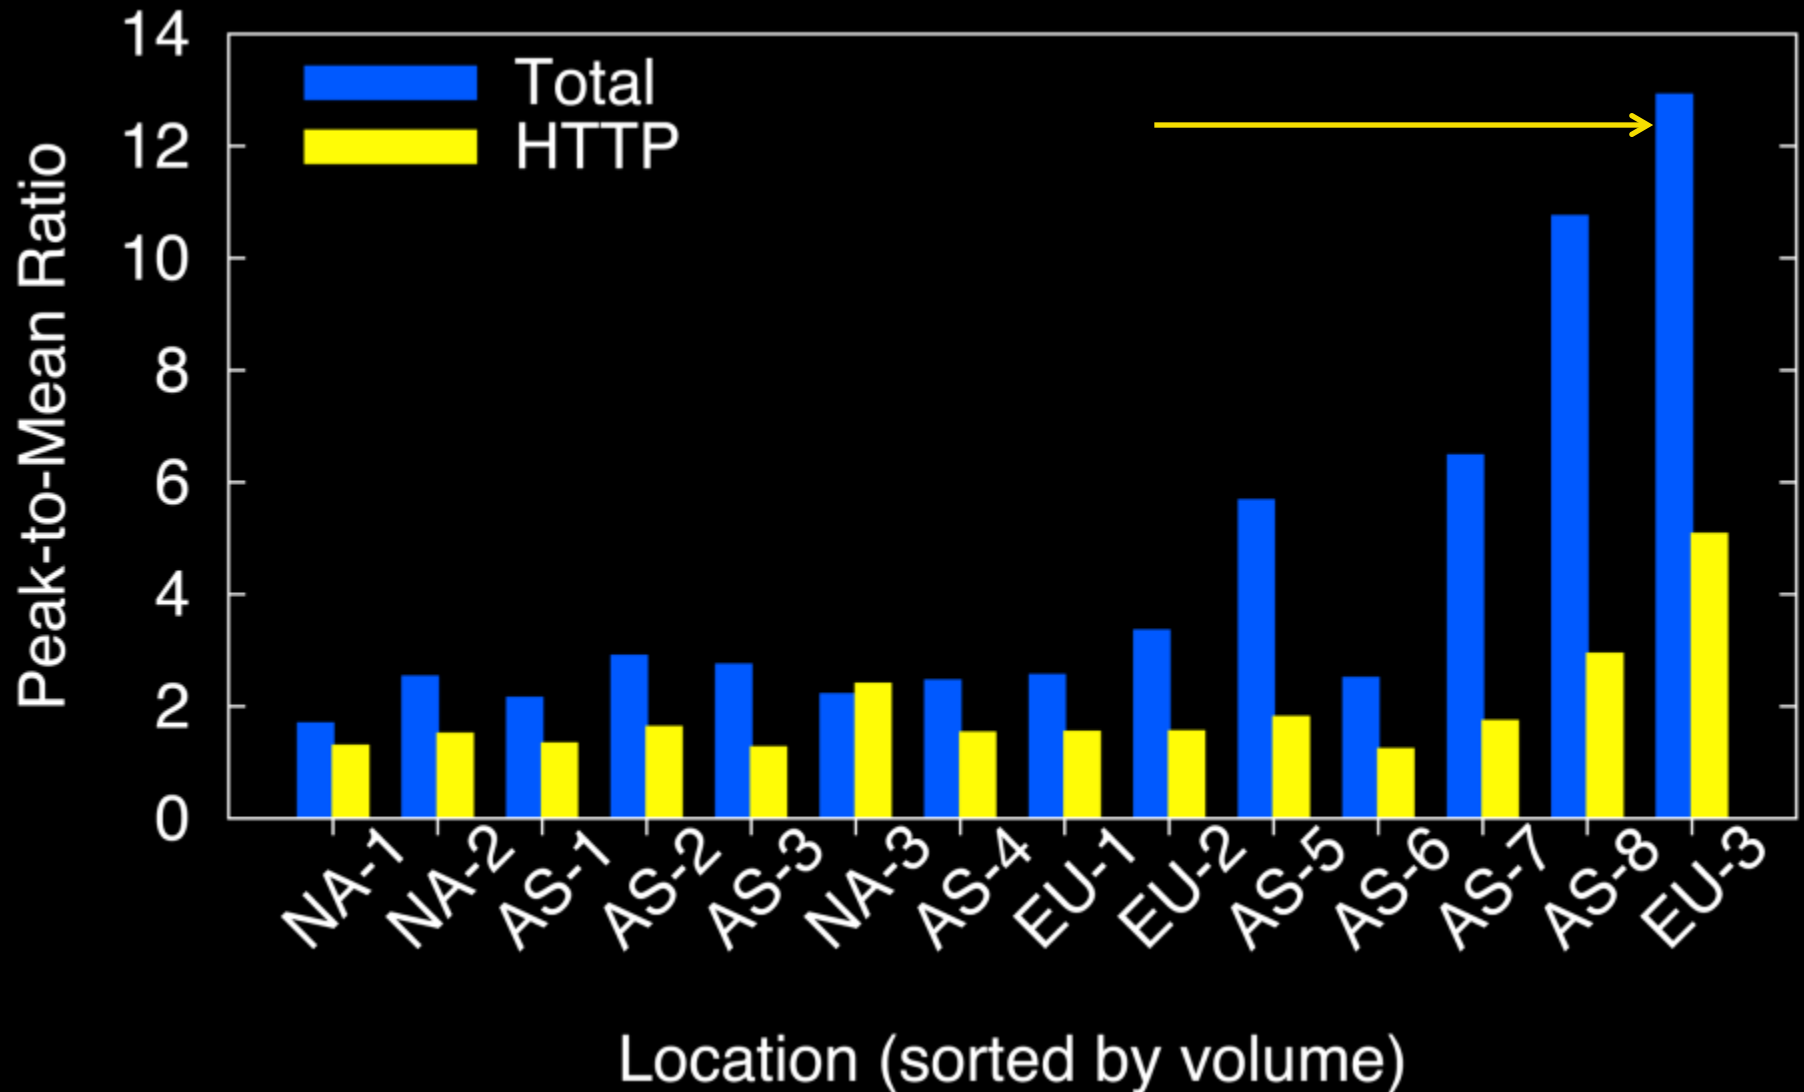
- **Bandwidth:** download times with BitTorrent increased on average 2.3%
- **Jitter:** consistently within industry standard bounds of 30ms

Does APLOMB negate the benefits  
of bandwidth-saving devices?



**APLOMB+ incurs a median penalty of 3.8% bandwidth inflation over traditional WAN Optimizers.**

Does “elastic scaling” at the cloud  
provide real benefits?



Some sites generate as much as 13x traffic more than average at peak hours.

# Recap

- Good application performance
  - Latency median inflation 1.1ms
  - Download times increased only 2.3%
- Generic redundancy elimination saves bandwidth costs
- Strong benefits from elasticity

# Conclusion

Moving middleboxes to the cloud is a **practical** and **feasible** solution to the complexity of enterprise networks.






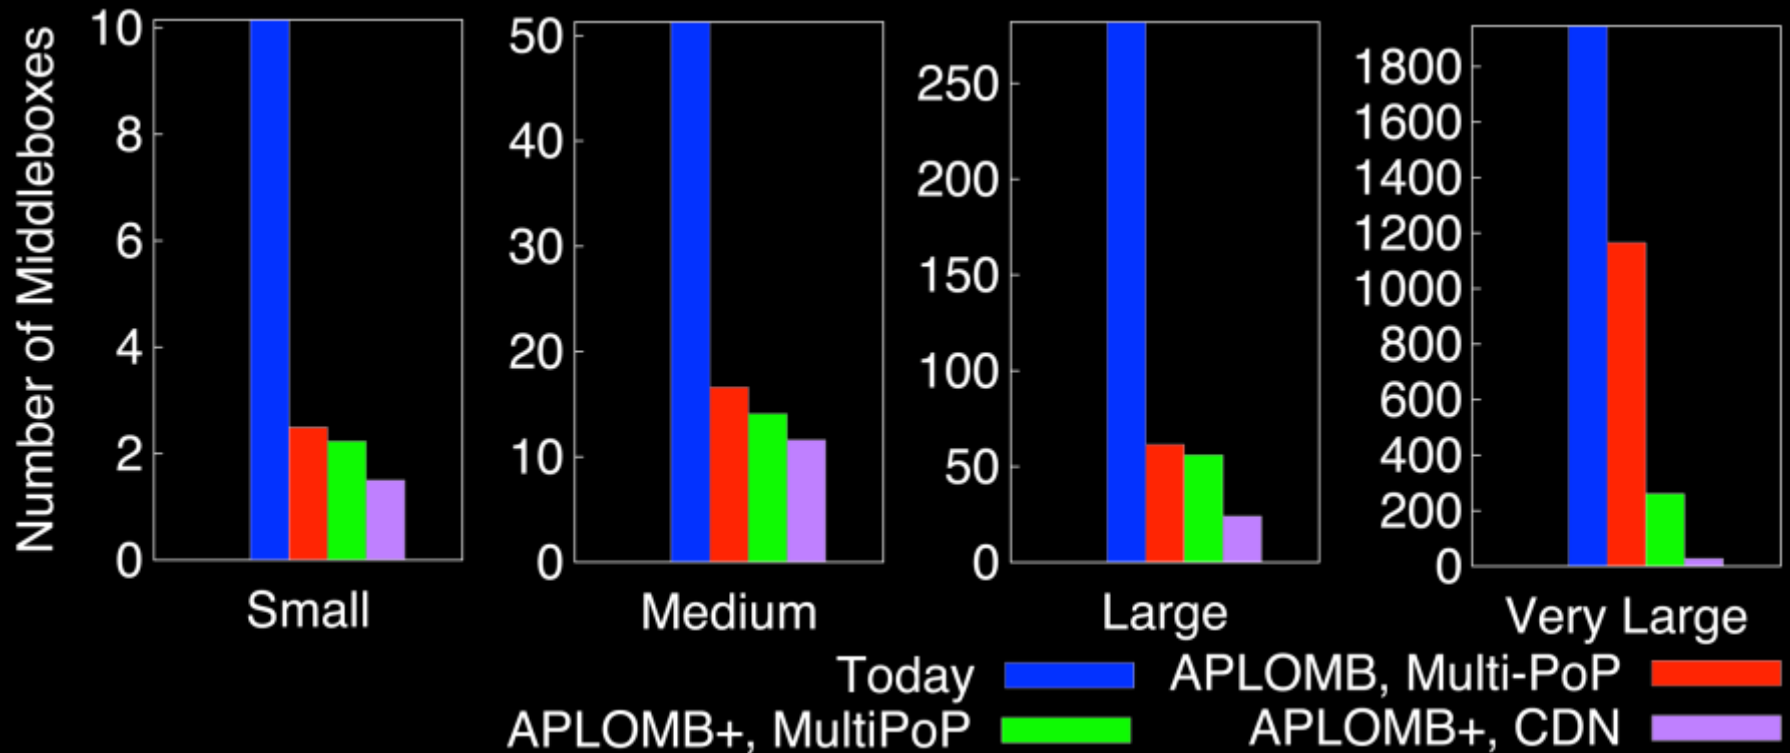
# What does it mean to “manage” middleboxes?

- Upgrades and Vendor Interaction
- Monitoring and Diagnostics
- Configuration
  - Appliance Configuration
  - Policy Configuration
- Training


# Internal Firewalls



# How many middleboxes can APLOMB outsource?

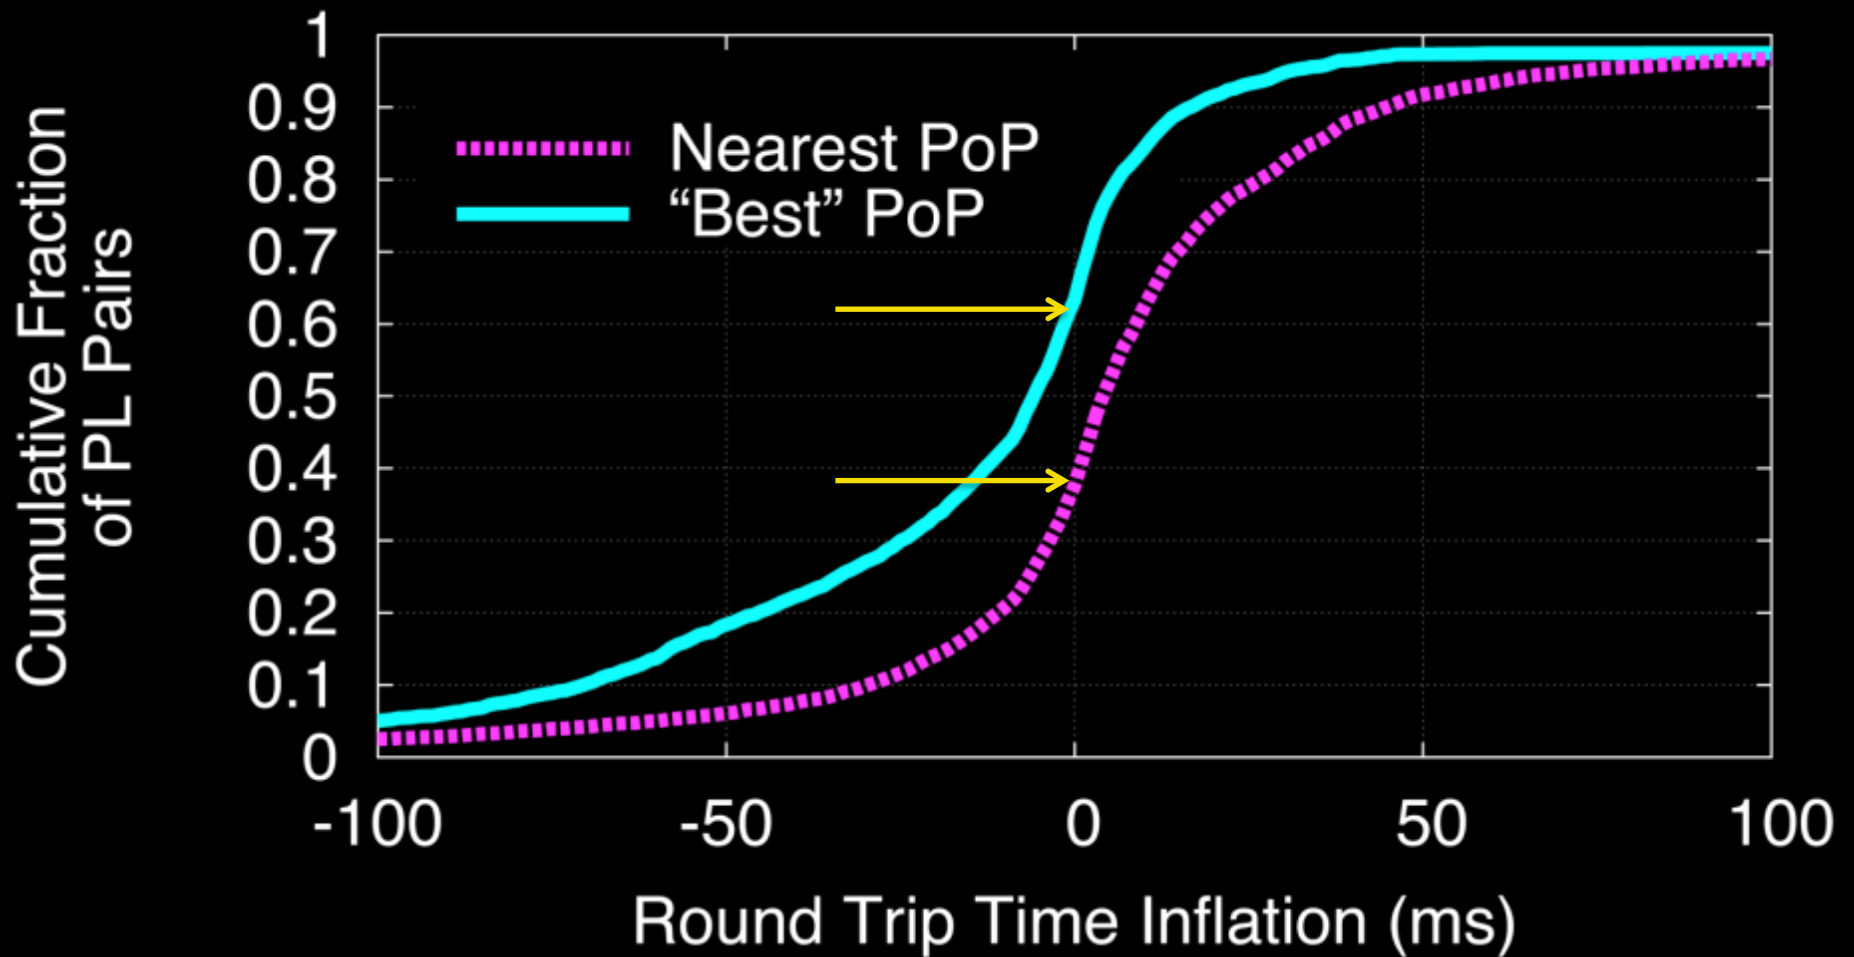


# How much do middleboxes cost?



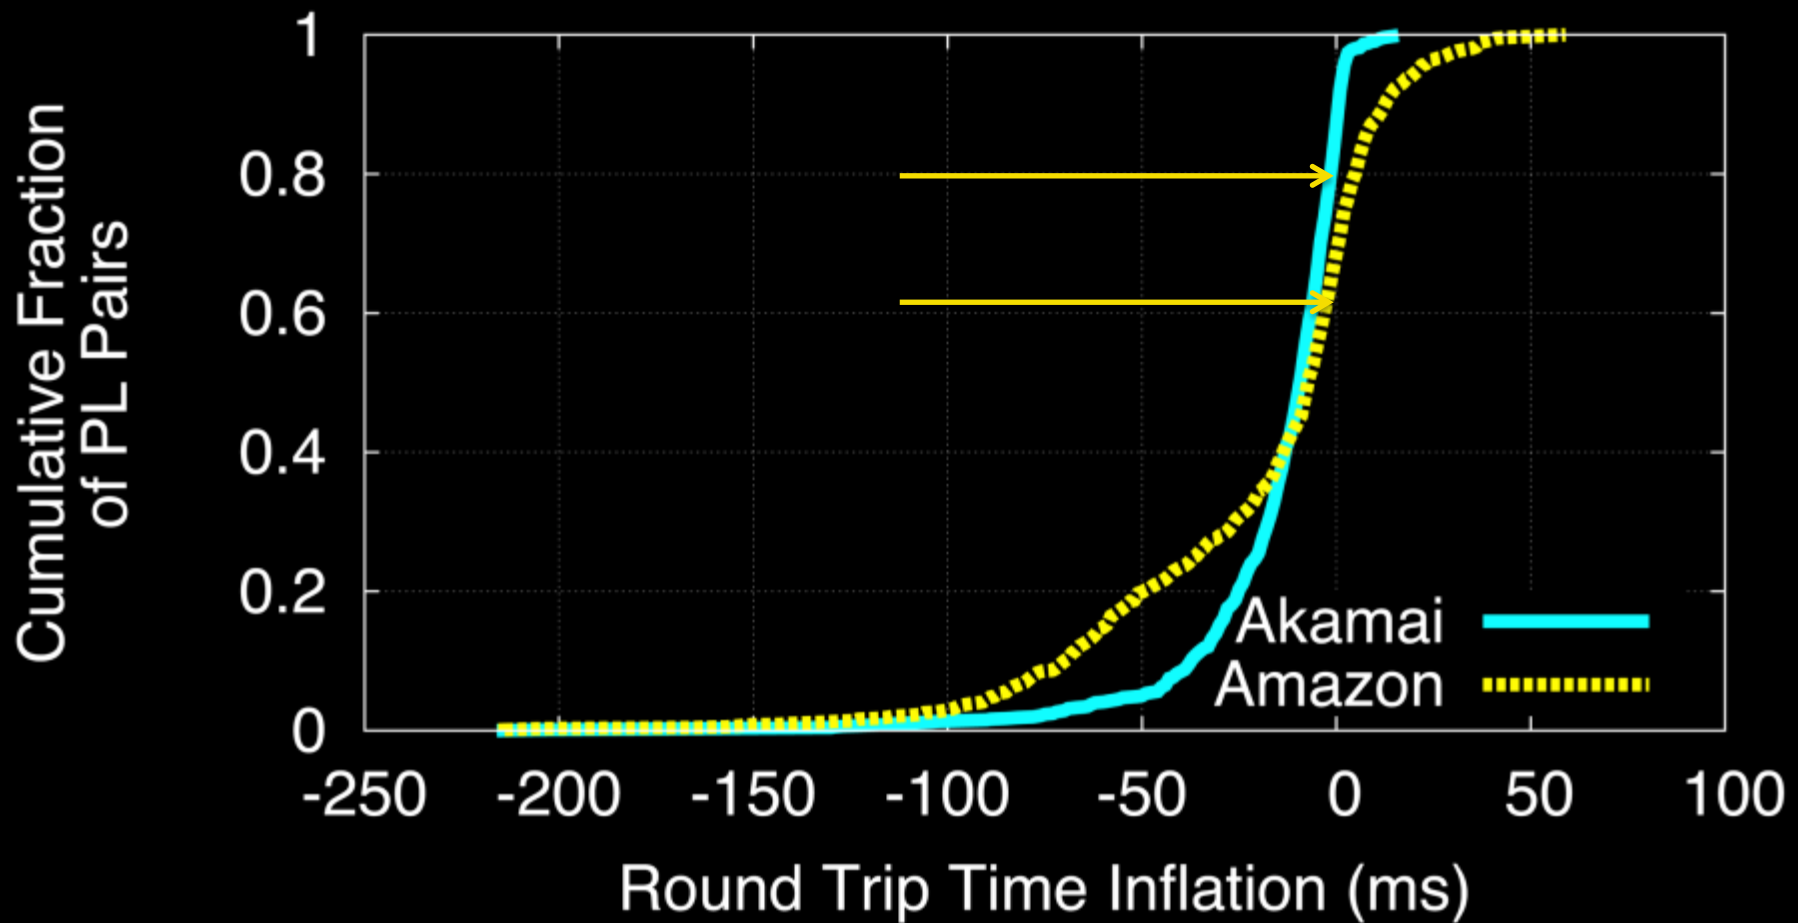
Thousands to millions of dollars / 5 years

Is maintaining multiple tunnels at the APLOMB gateway useful?



With multiple tunnels, the fraction of pairs with 0 inflation or better moves from 40% to 60%


How large must a provider's  
datacenter footprint be to support  
middlebox services?



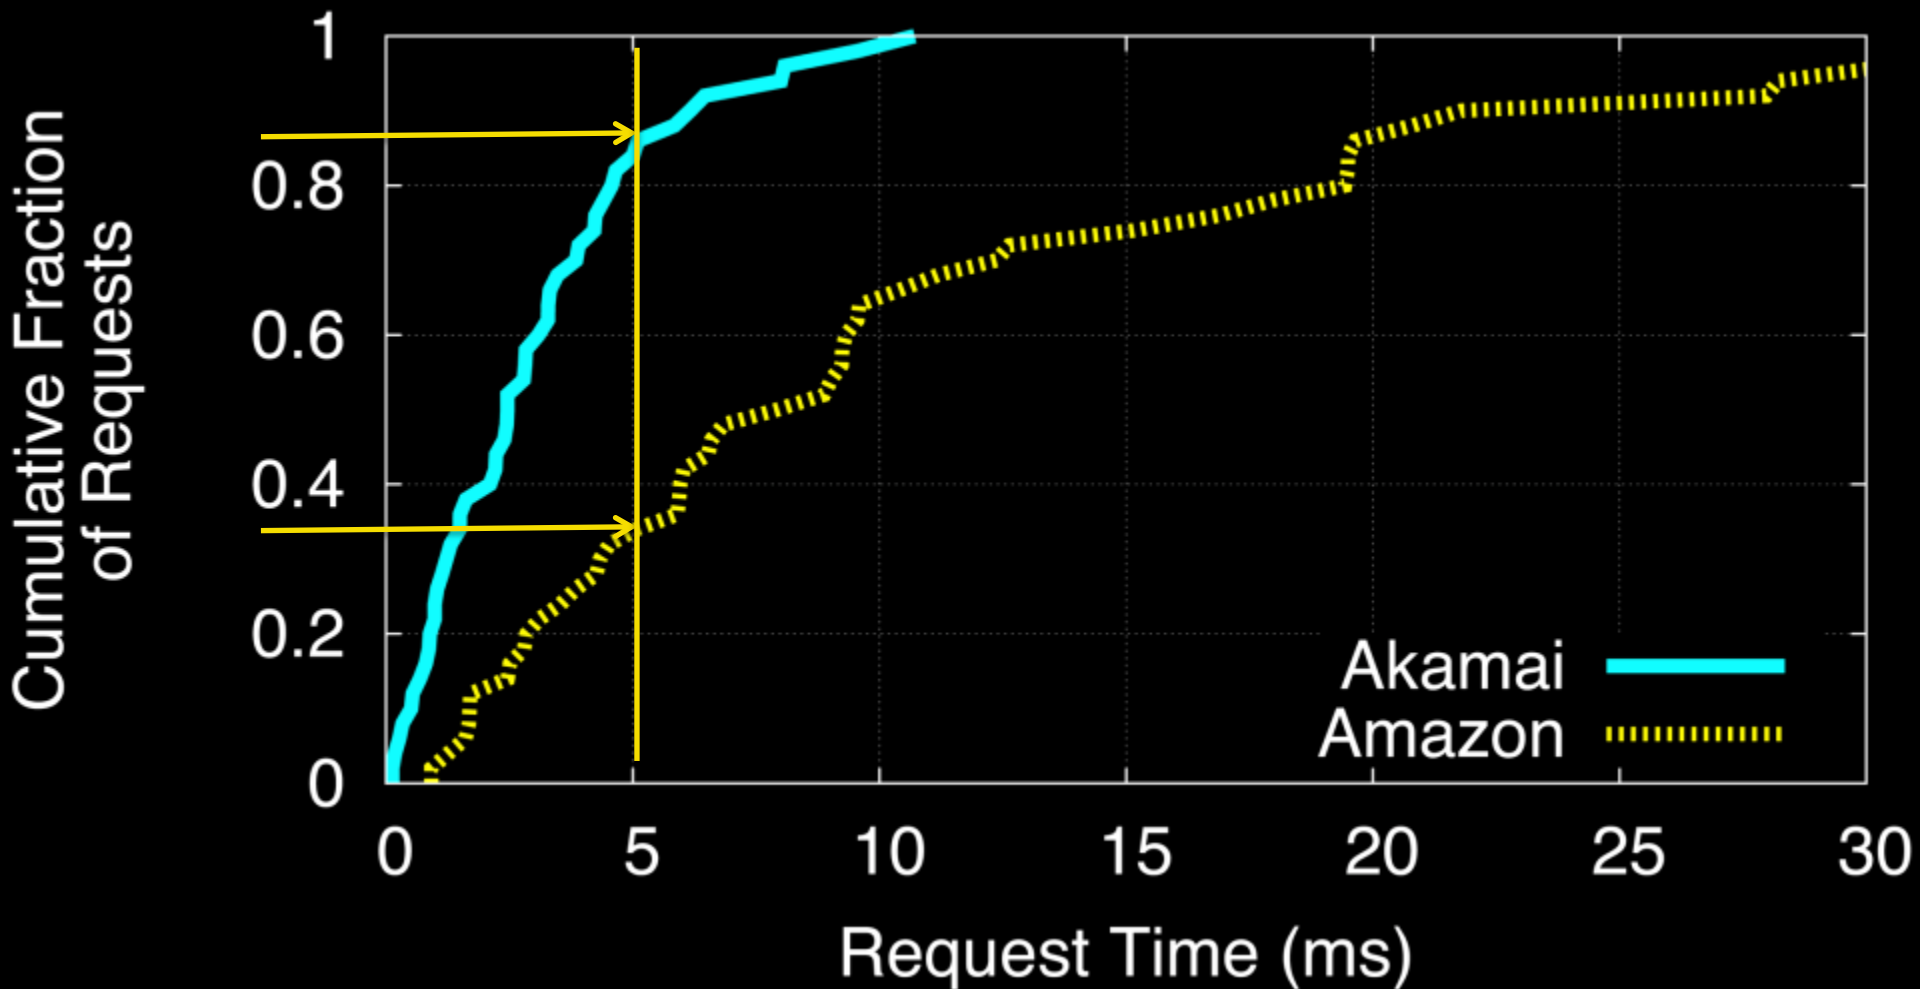
**Minimal Improvement to E2E Latency with  
Larger Footprint.**



How does APLOMB redirection  
impact web page load times?



Median: slightly worse; 90%-ile: slightly better.



Caches may require a larger footprint to provide nationwide service.

