

Data Center Middleboxes

Hakim Weatherspoon

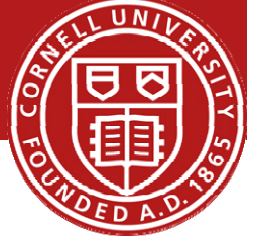
Assistant Professor, Dept of Computer Science

CS 5413: High Performance Systems and Networking

November 24, 2014

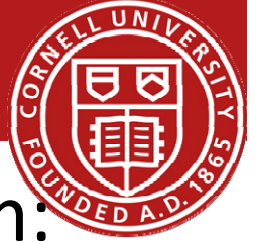
Slides from ACM SIGCOMM 2012 presentation on “Making middleboxes someone else's problem: network processing as a cloud service”

Where are we in the semester?



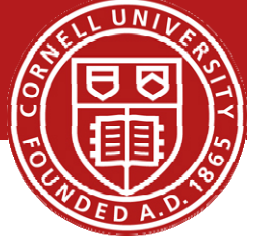
- Overview and Basics
- Data Center Networks
 - Basic switching technologies
 - Data Center Network Topologies (today and Monday)
 - Software Routers (eg. Click, Routebricks, NetMap, Netslice)
 - Alternative Switching Technologies
 - Data Center Transport
- Data Center Software Networking
 - Software Defined networking (overview, control plane, data plane, NetFGPA)
 - Data Center Traffic and Measurements
 - Virtualizing Networks
 - Middleboxes
- Advanced Topics

Goals for Today



- Making middleboxes someone else's problem: network processing as a cloud service,
 - J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar. ACM SIGCOMM Computer Communication Review (CCR) Volume 42, Issue 4 (August 2012), pages 13-24.

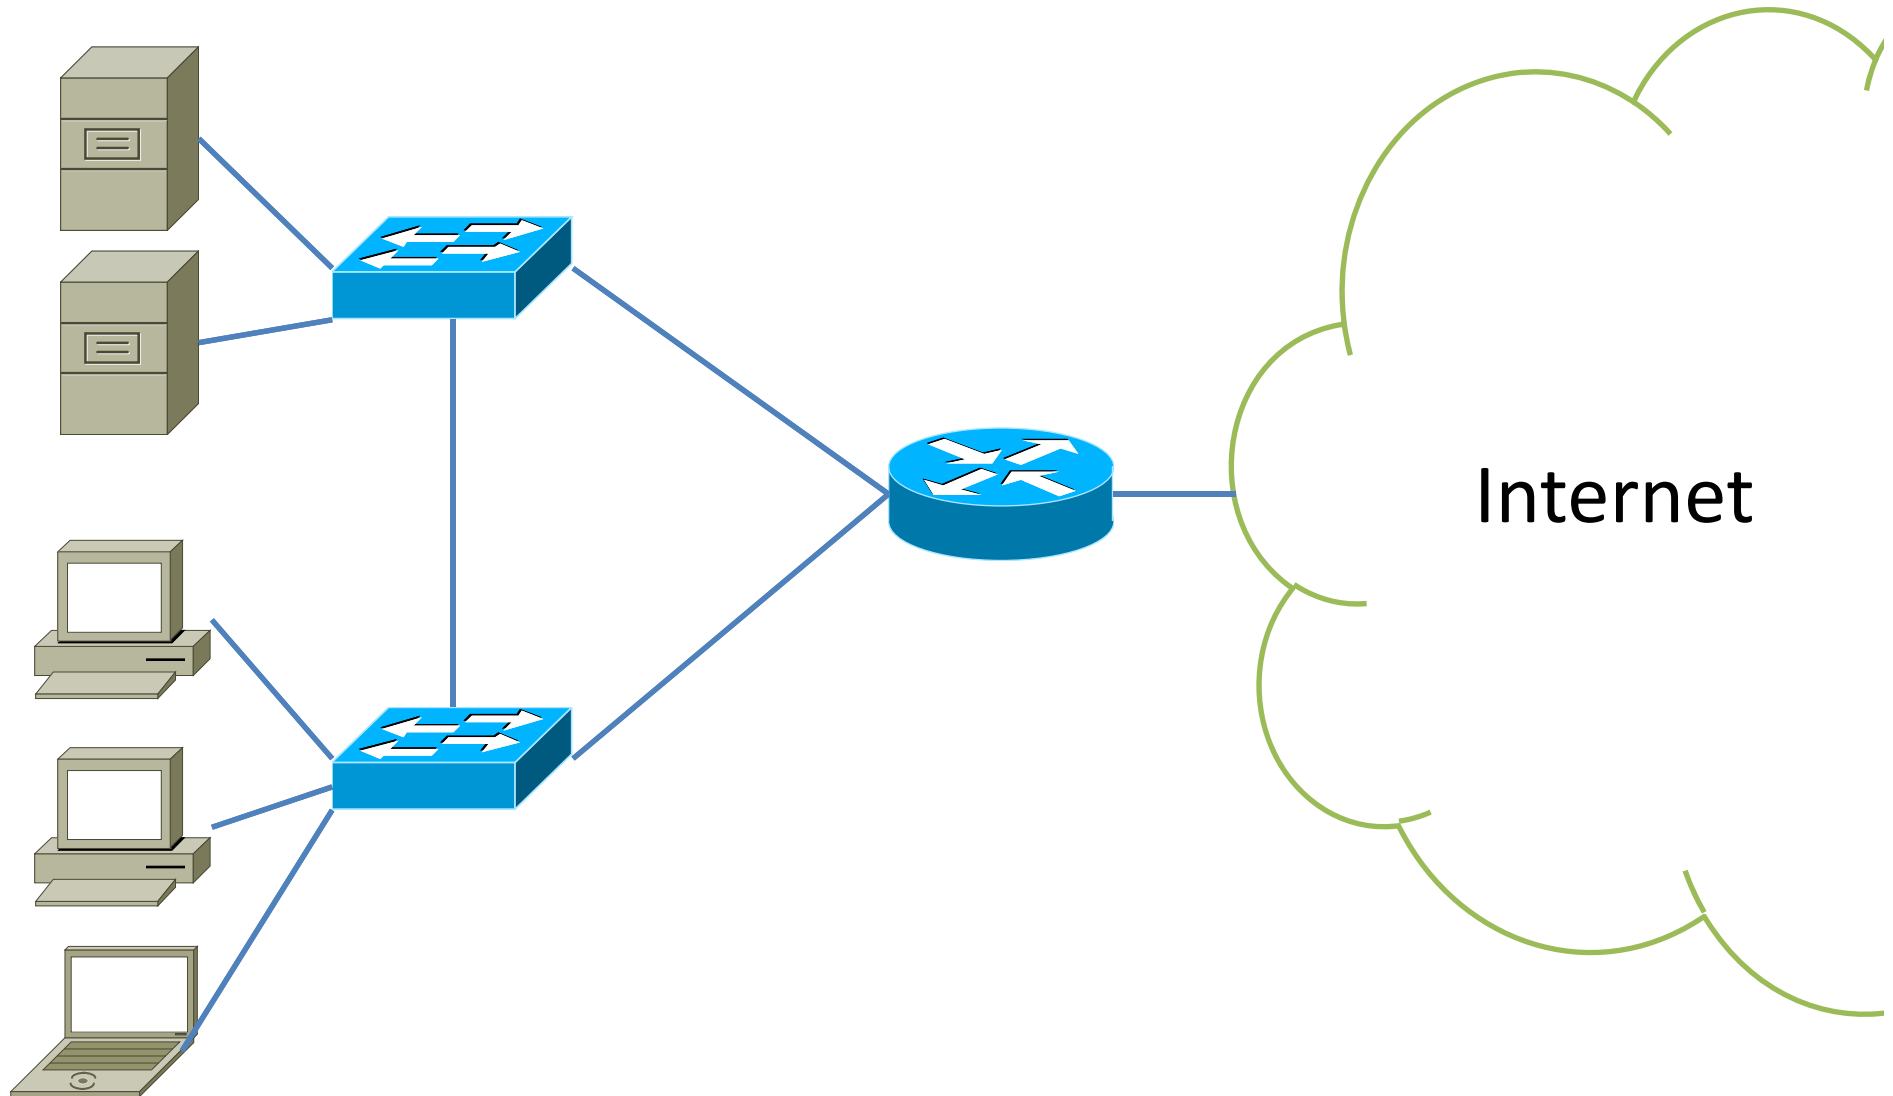
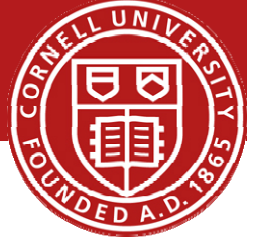
APLOMB



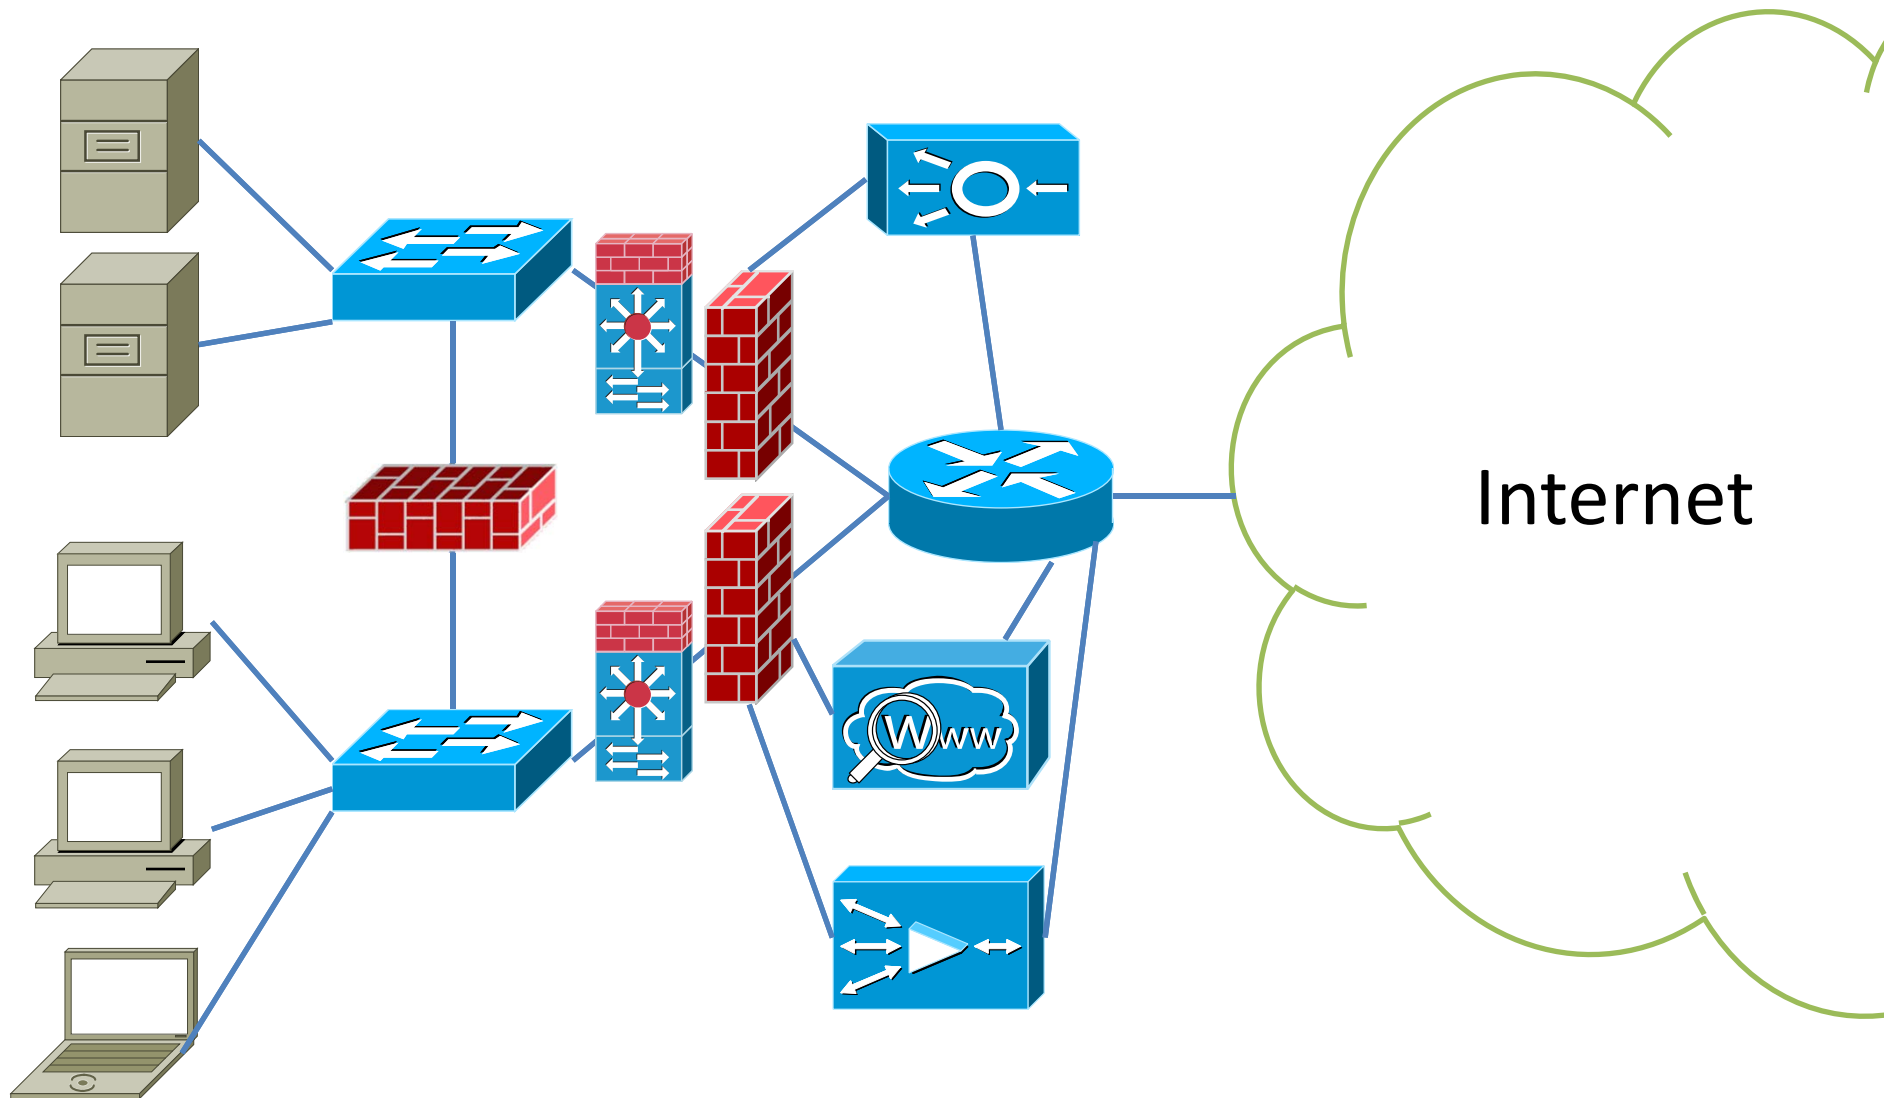
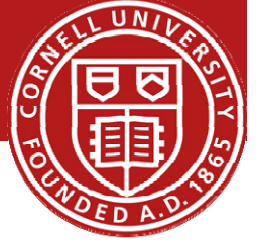
“Appliance for Outsourcing Middleboxes”

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

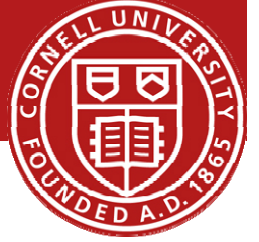
Typical Enterprise Networks



Typical Enterprise Networks

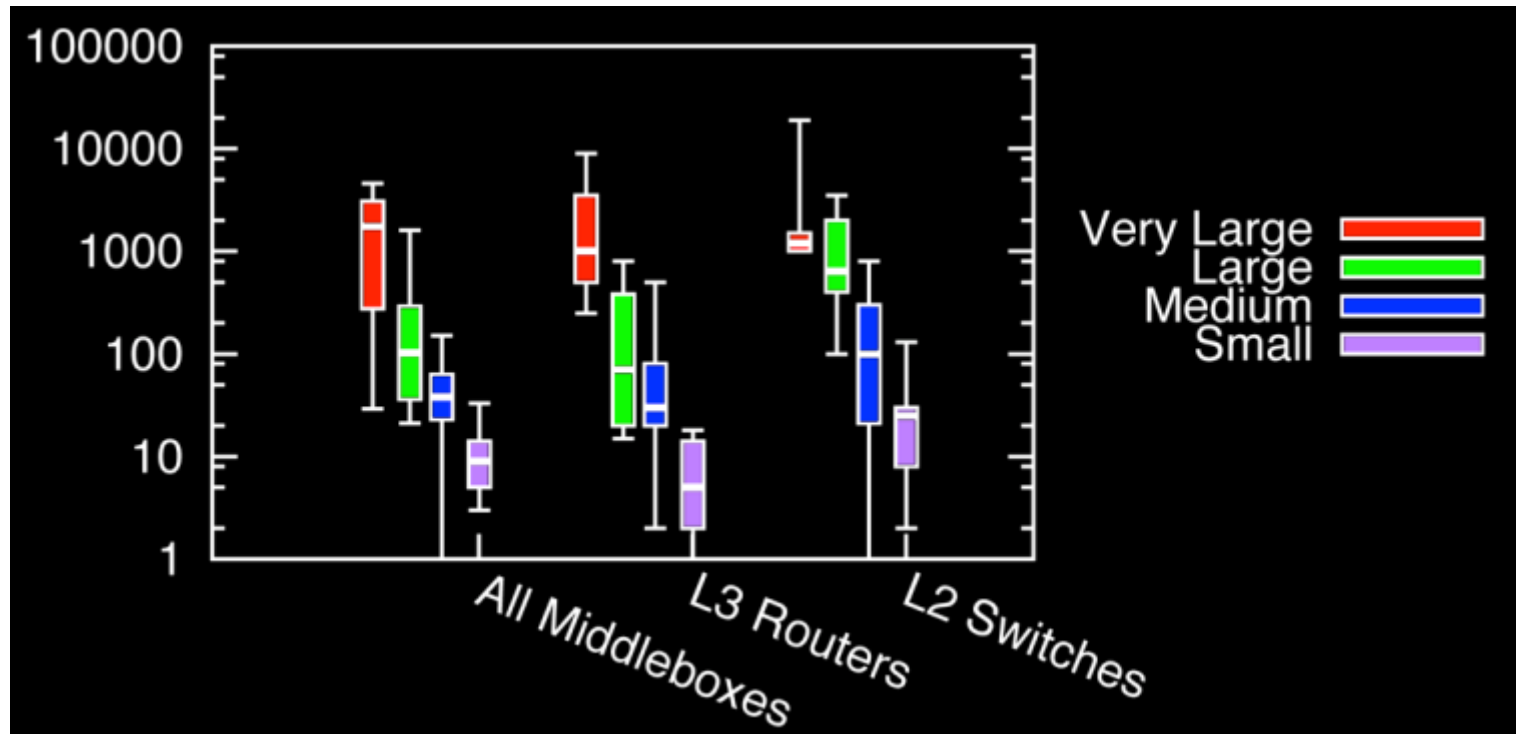
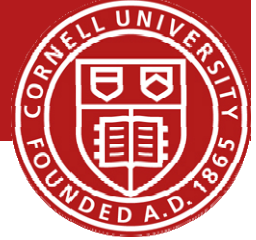


A Survey



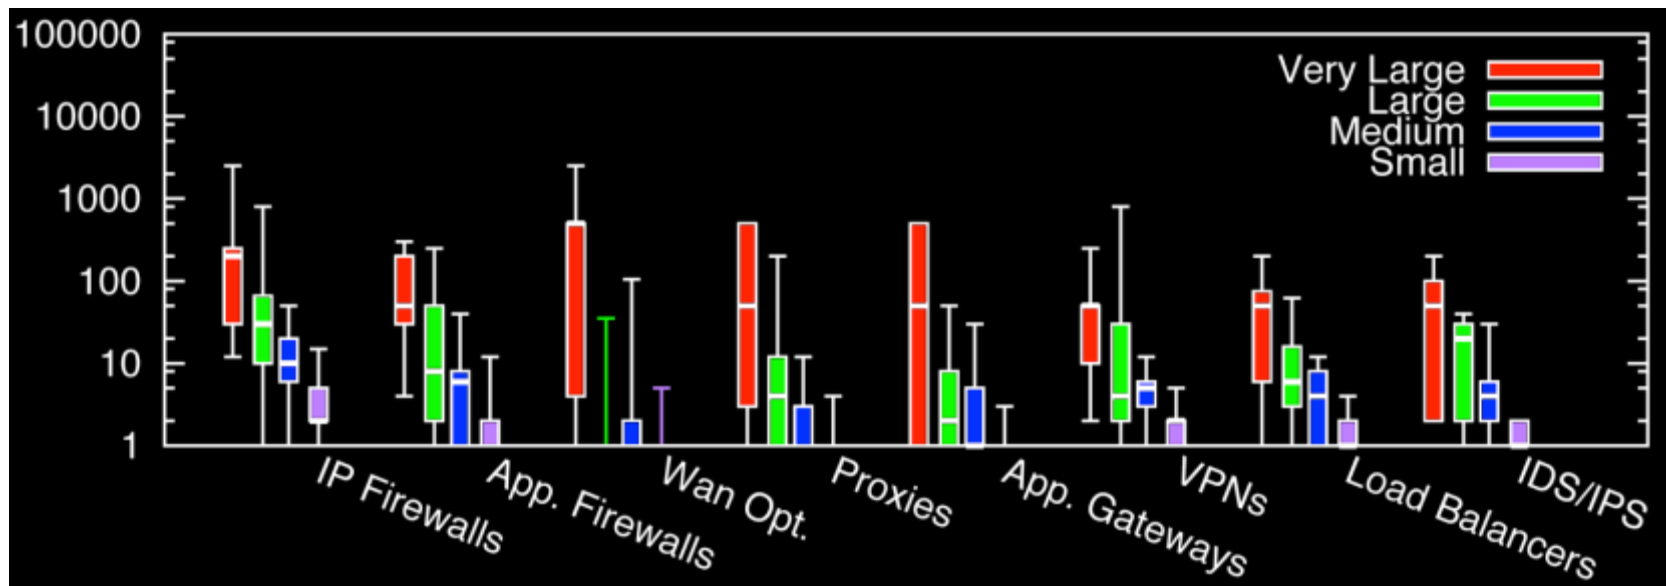
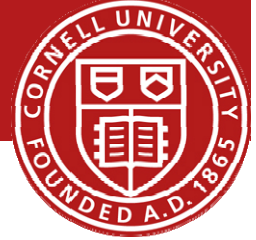
- 57 enterprise network administrators
- Small (< 1k hosts) to XL (>100k hosts)
- Asked about deployment size, expenses, complexity, and failures.

How many middleboxes do you deploy?



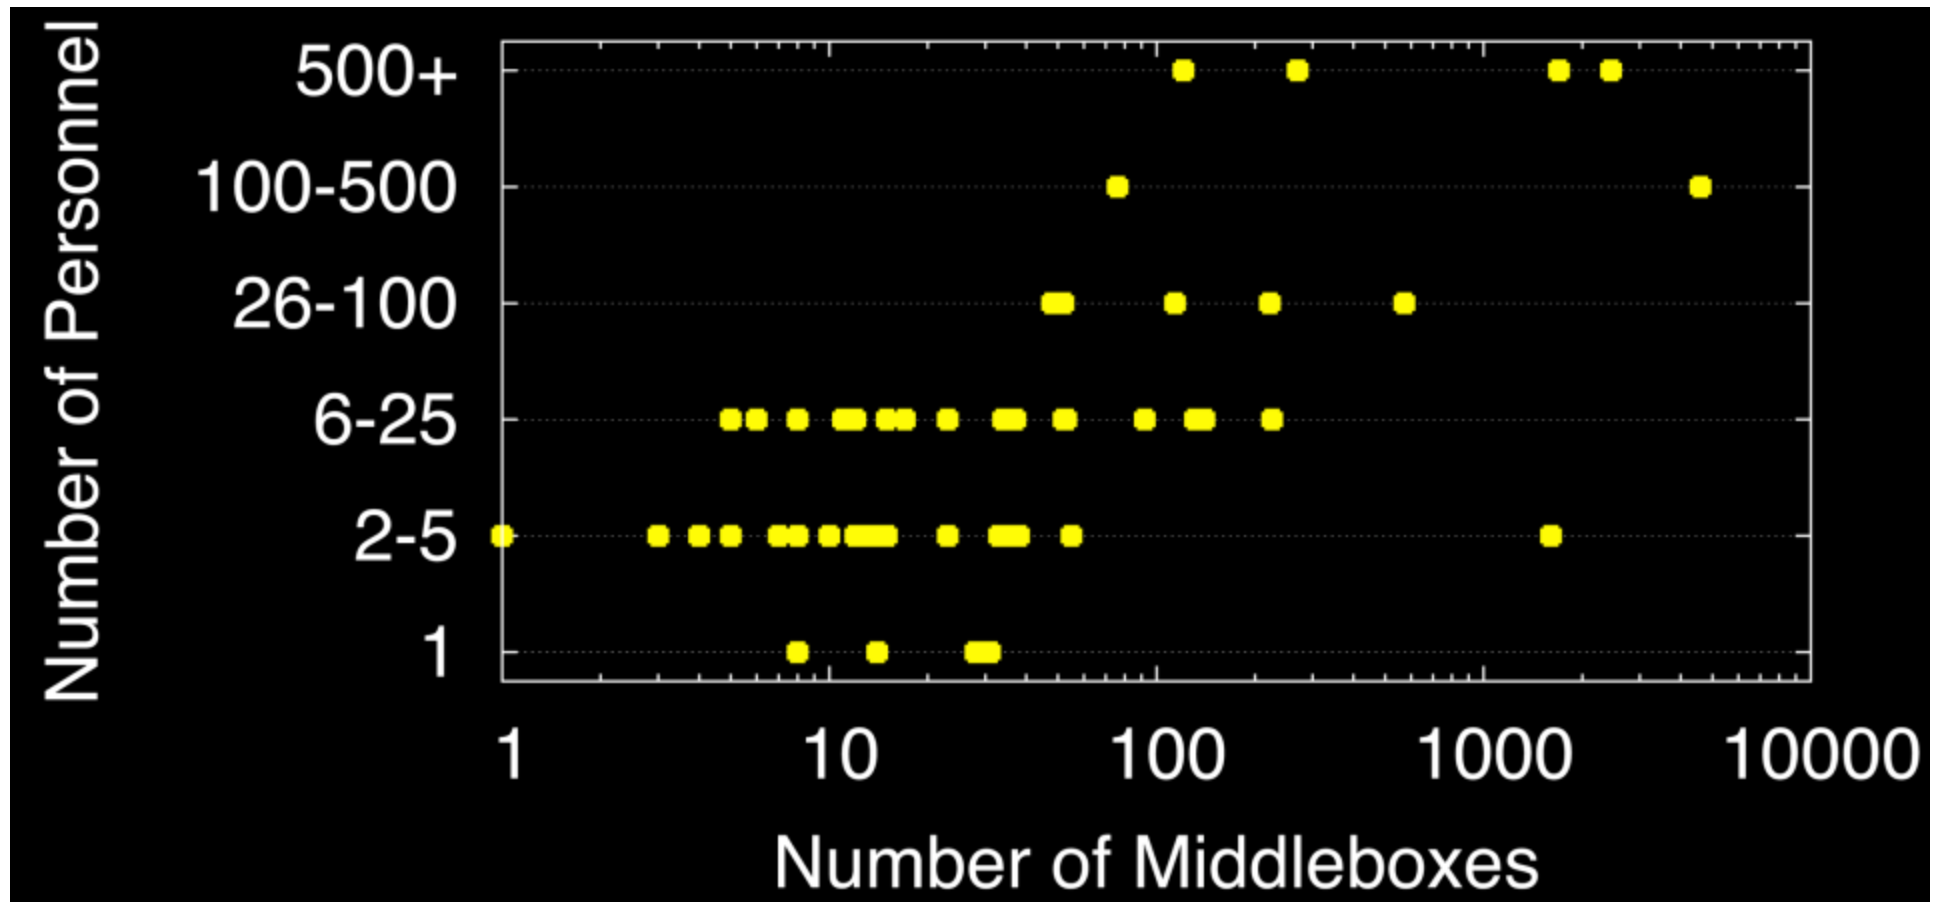
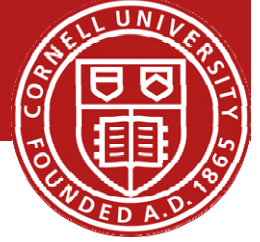
Typically on par with # routers and switches.

What kinds of middleboxes do you deploy?

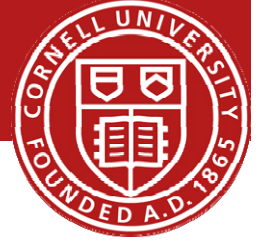


Many kinds of devices, all with different functions and management expertise required.

How many networking personnel are there?



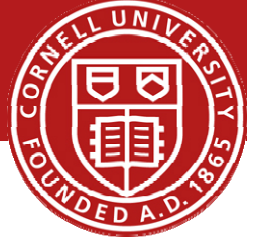
How do administrators spend their time?



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

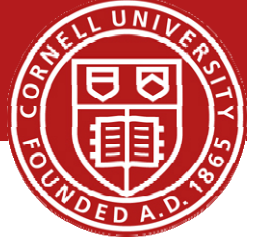
	Misconfig.	Overload	Physical/ Electrical
Firewalls	67.3%	16.3%	16.3%
Proxies	63.2%	15.7%	21.1%
IDS	54.45%	11.4%	34%

Recap

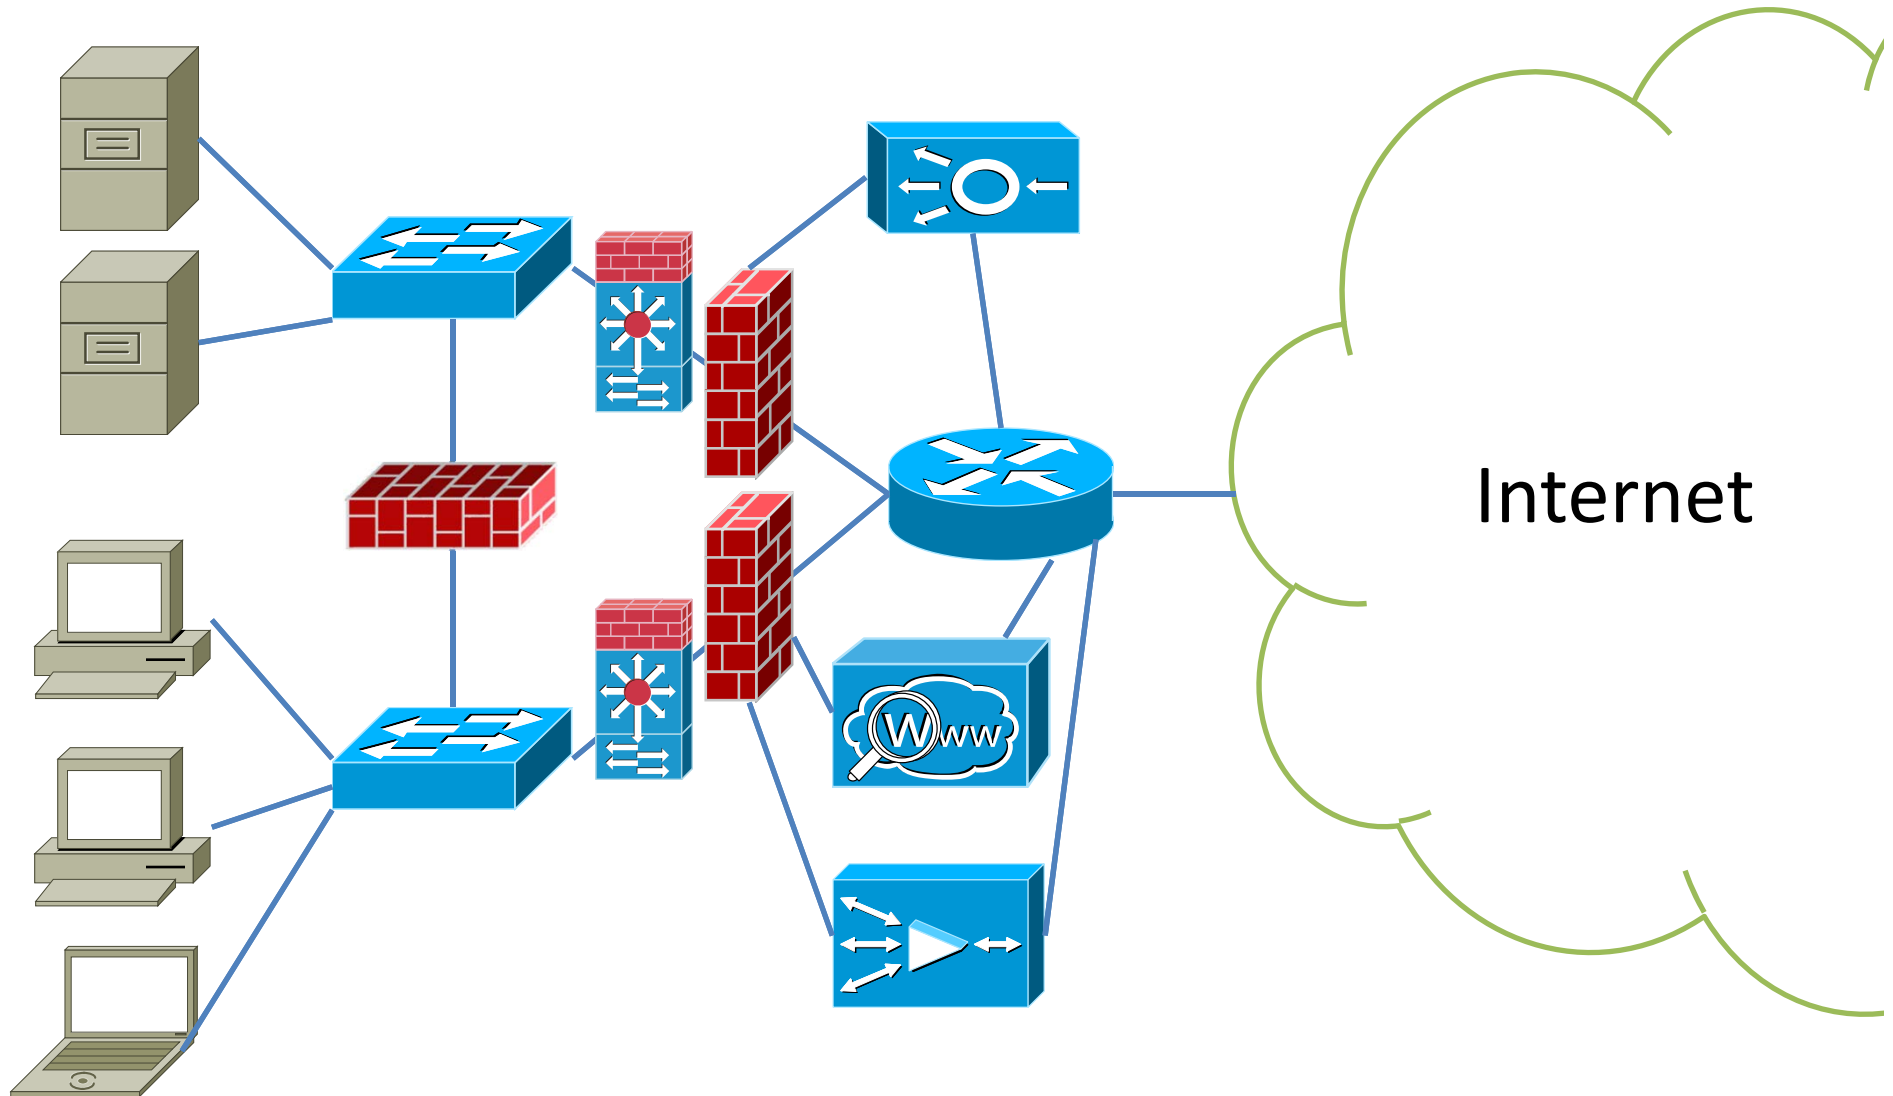
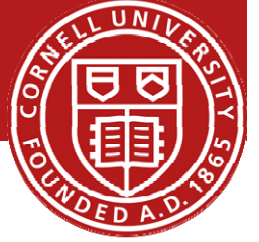


- High Capital and Operating Expenses
- Time Consuming and Error-Prone
- Physical and Overload Failures

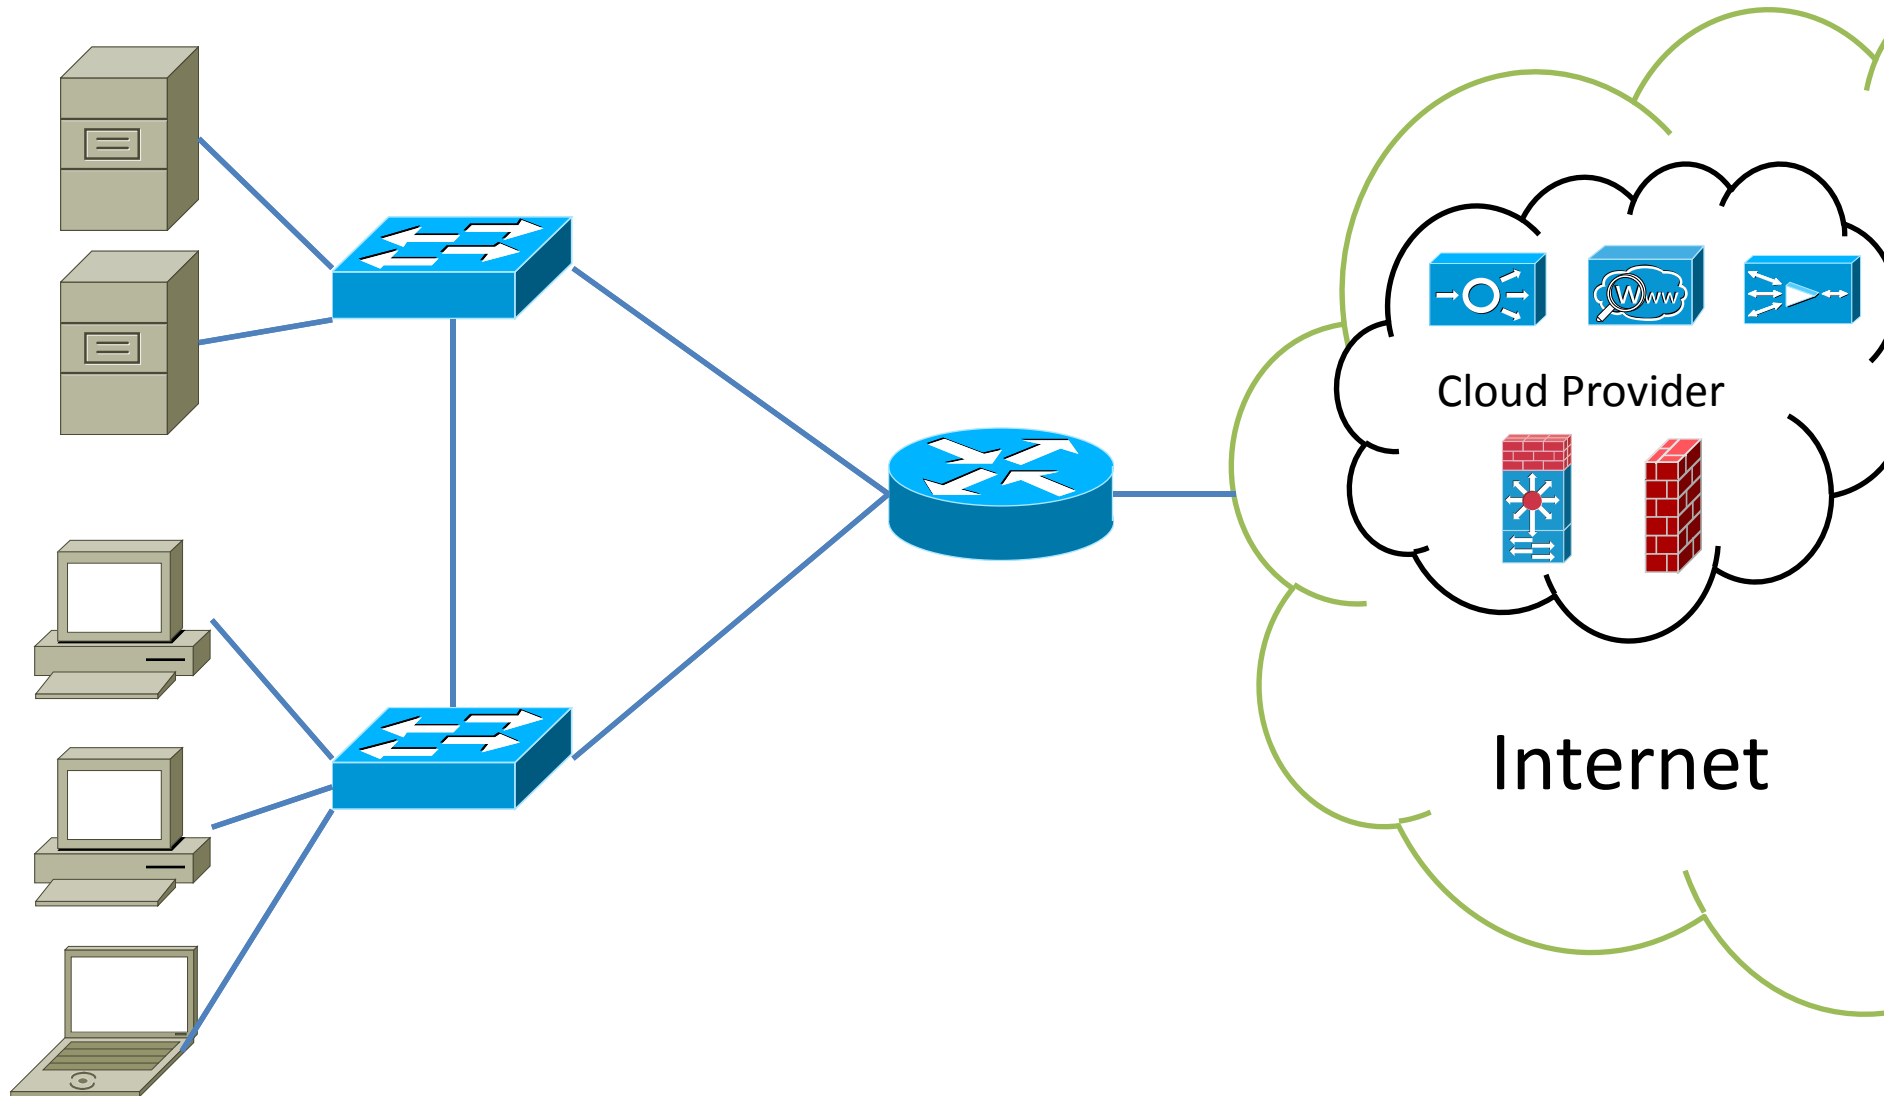
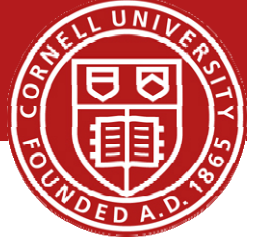
How can we improve this?



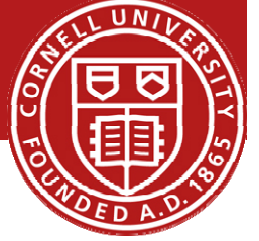
Proposal



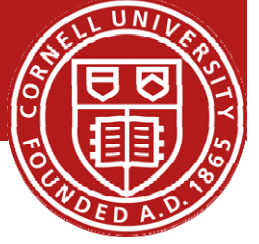
Proposal



A move to the cloud

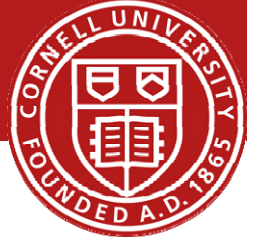


- High Capital and Operating Expenses
 - Economies of scale and pay-per use
- Time Consuming and Error Prone
 - Simplifies configuration and deployment
- Physical and Overload Failures
 - Redundant resources for failover

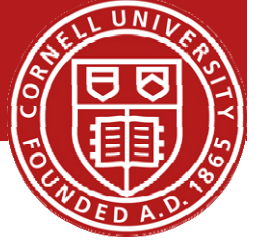


Design

Challenges



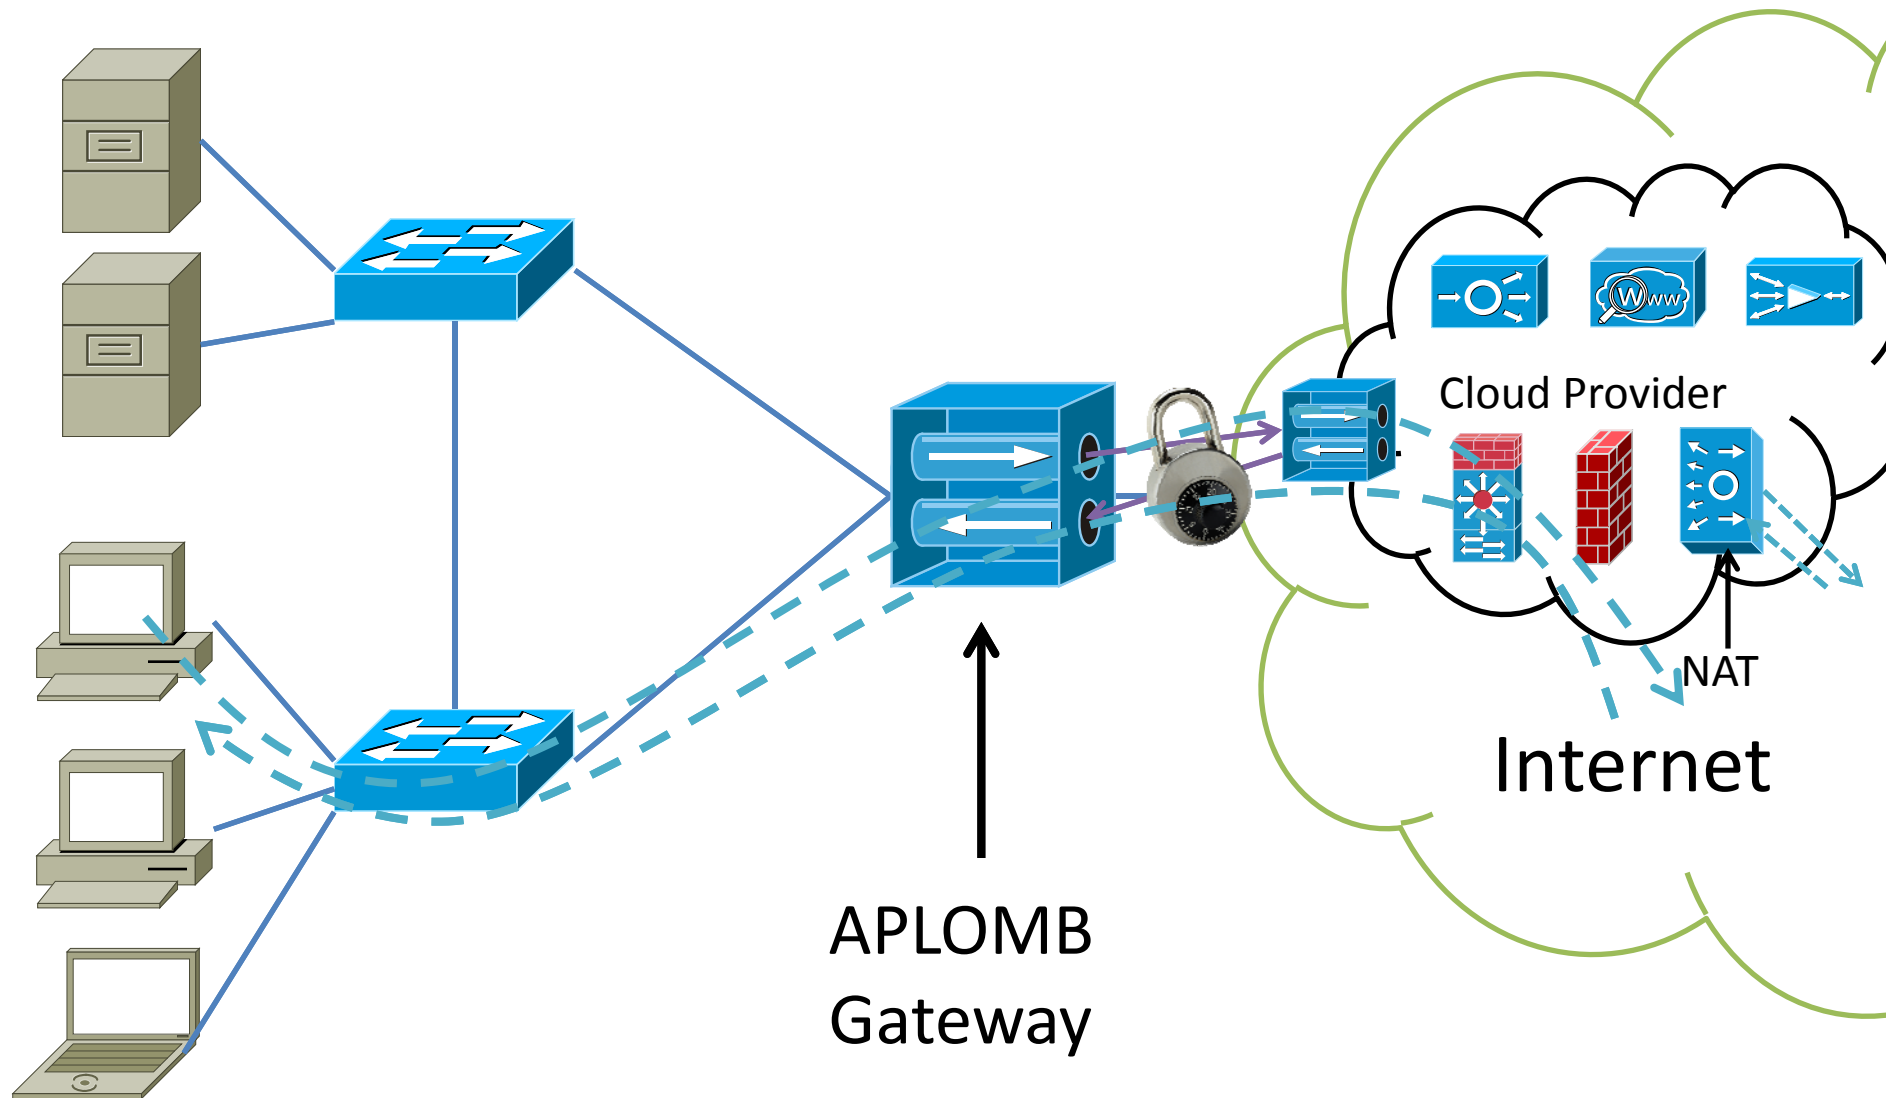
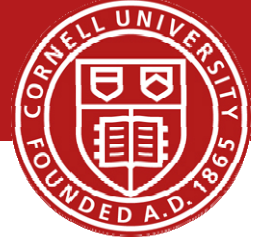
- Minimal Complexity at the Enterprise
- Functional Equivalence
- Low Performance Overhead



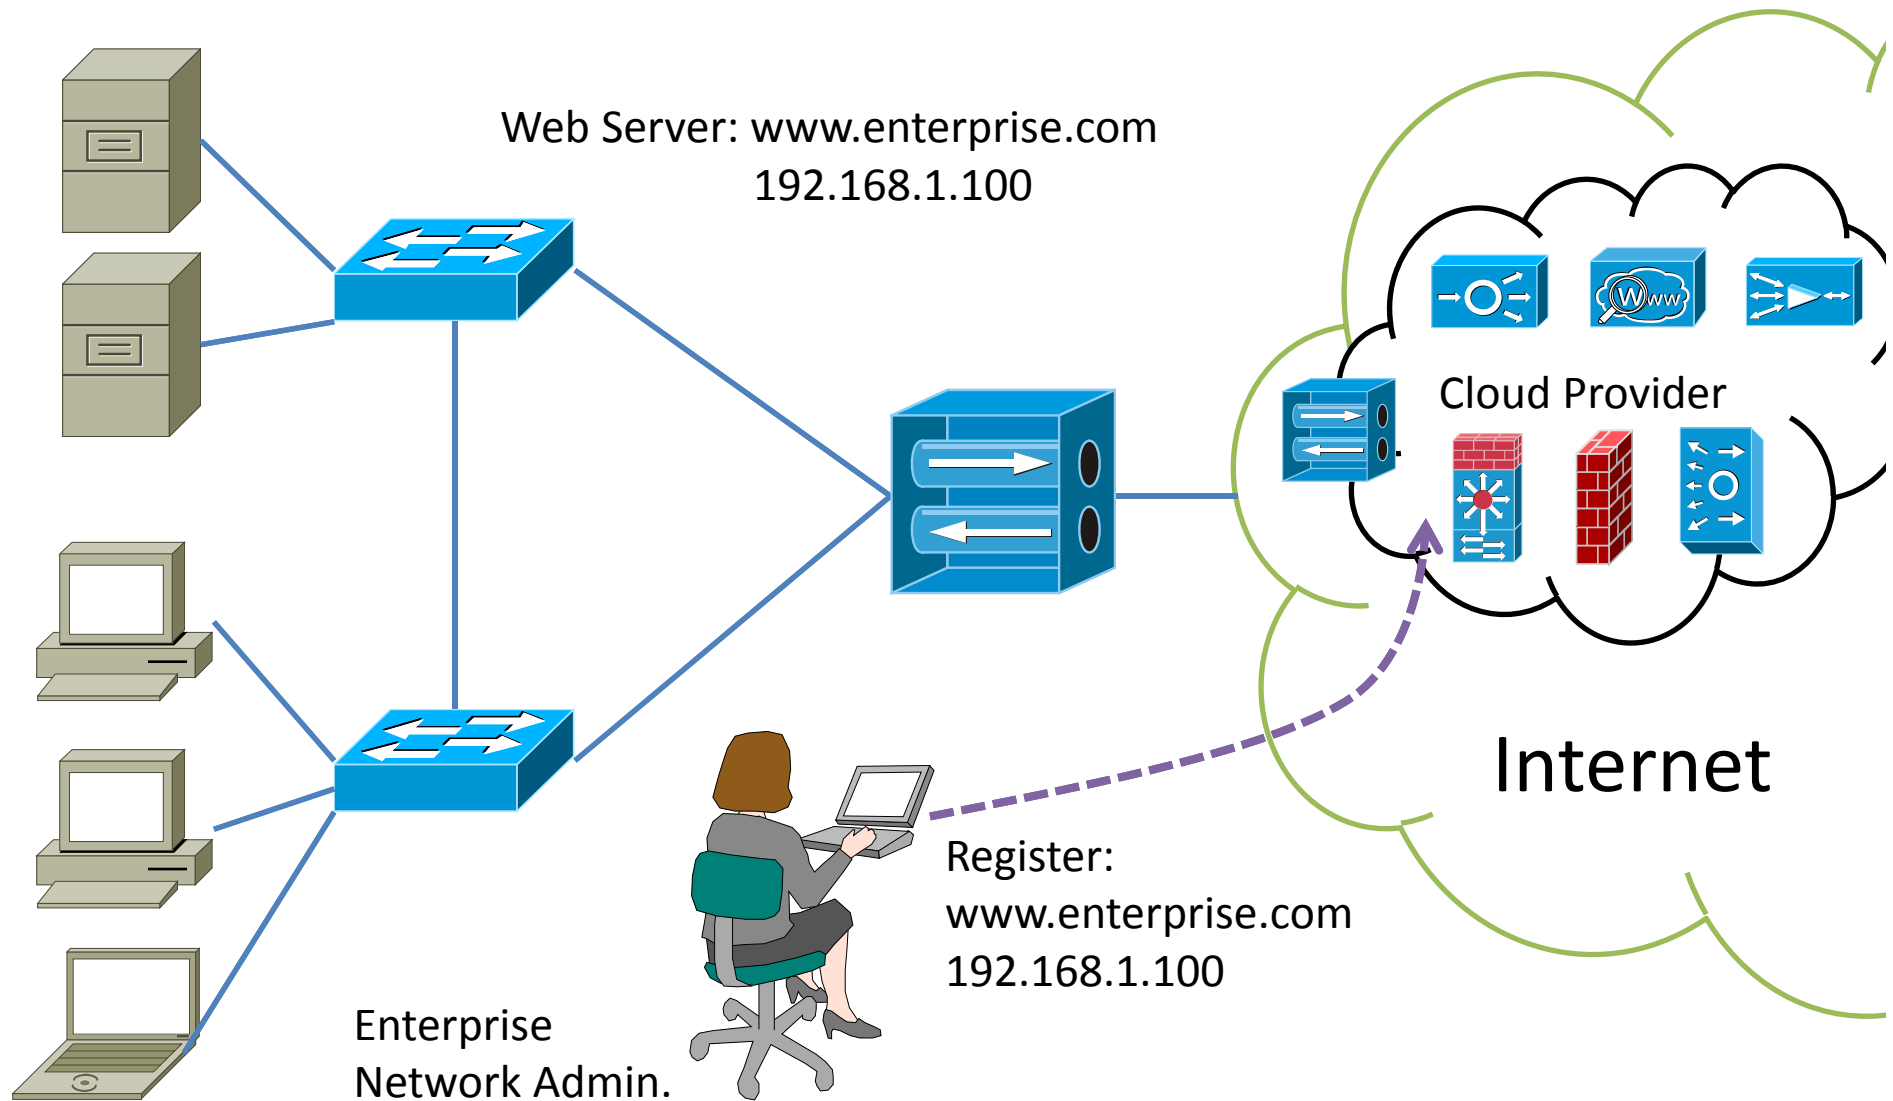
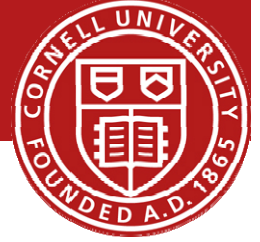
APLOMB

“Appliance for Outsourcing Middleboxes”

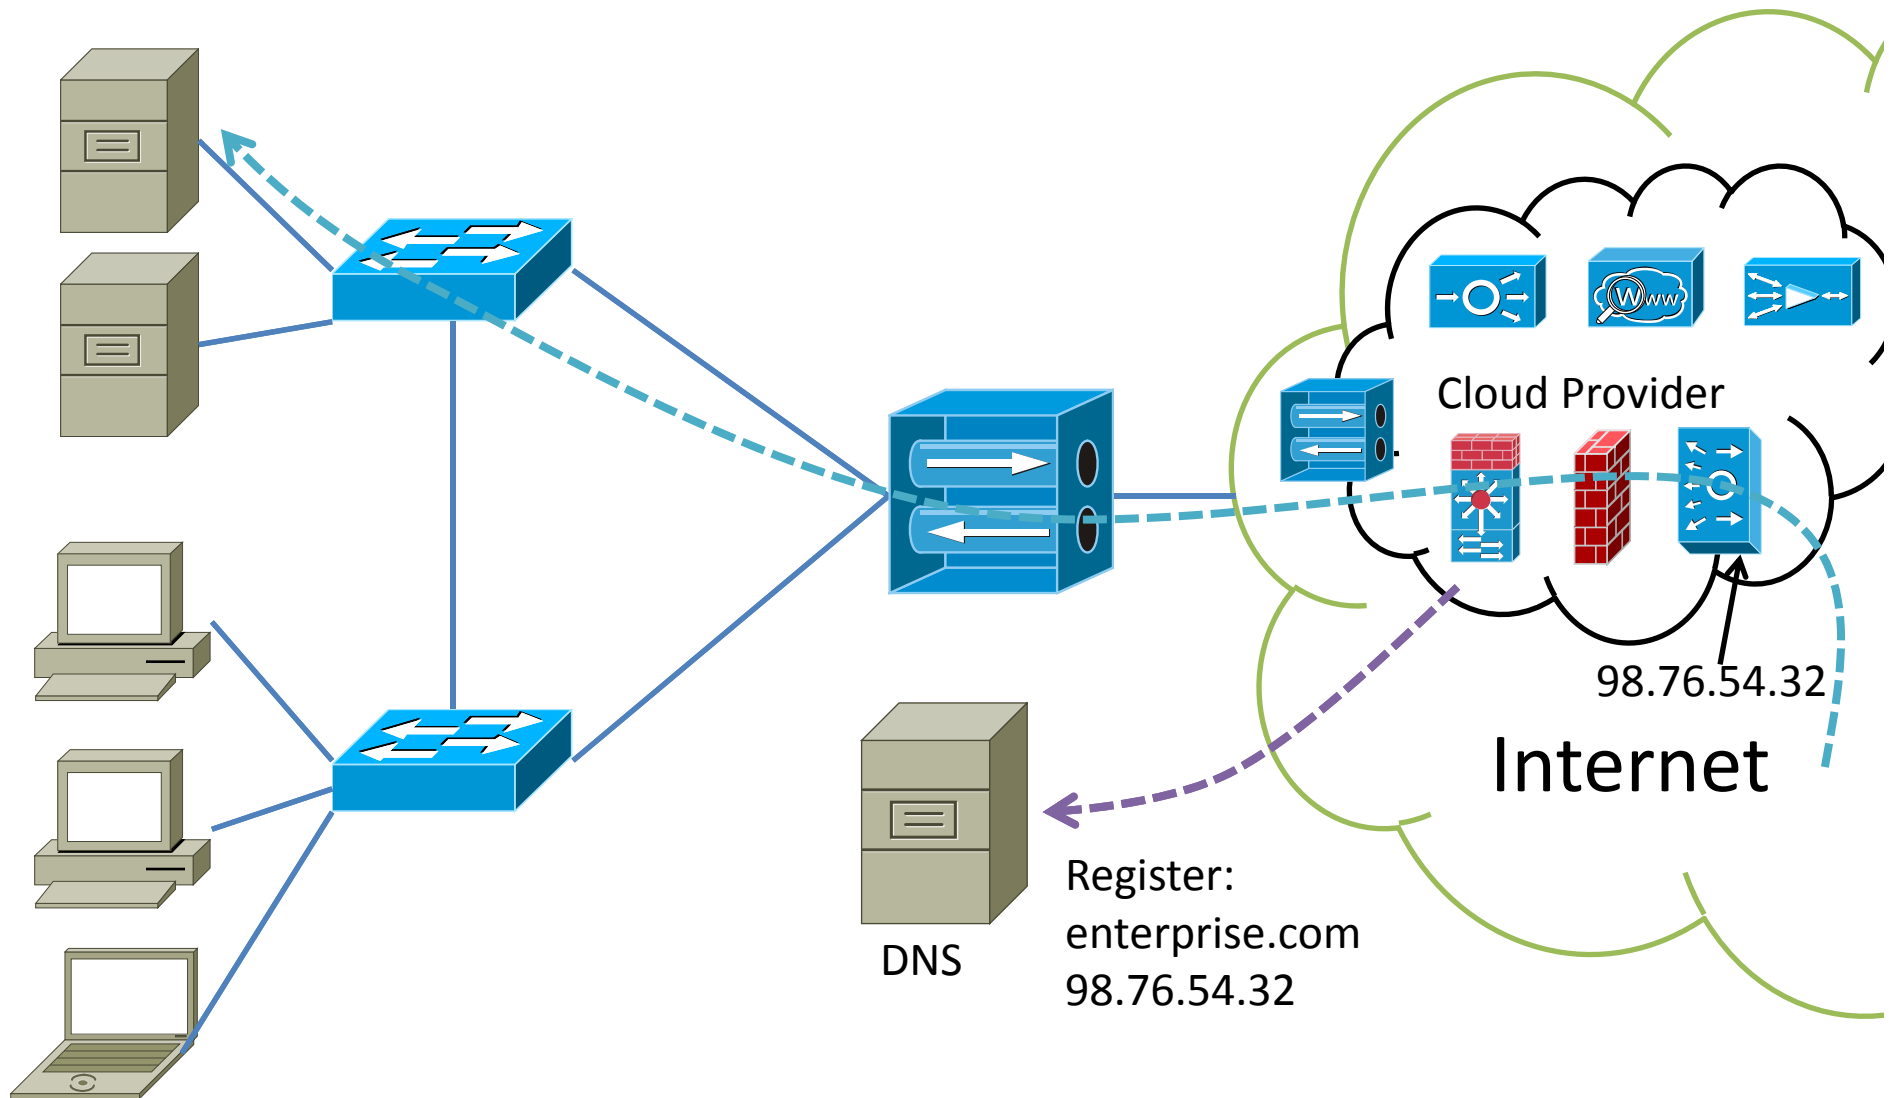
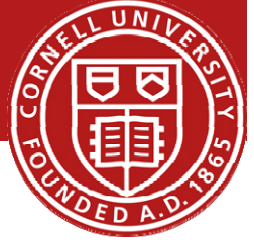
Outsourcing Middleboxes with APLOMB



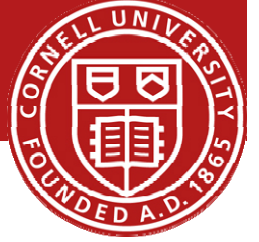
Inbound Traffic



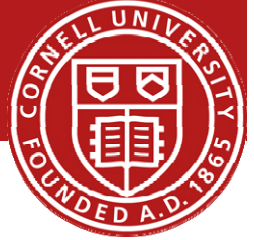
Inbound Traffic



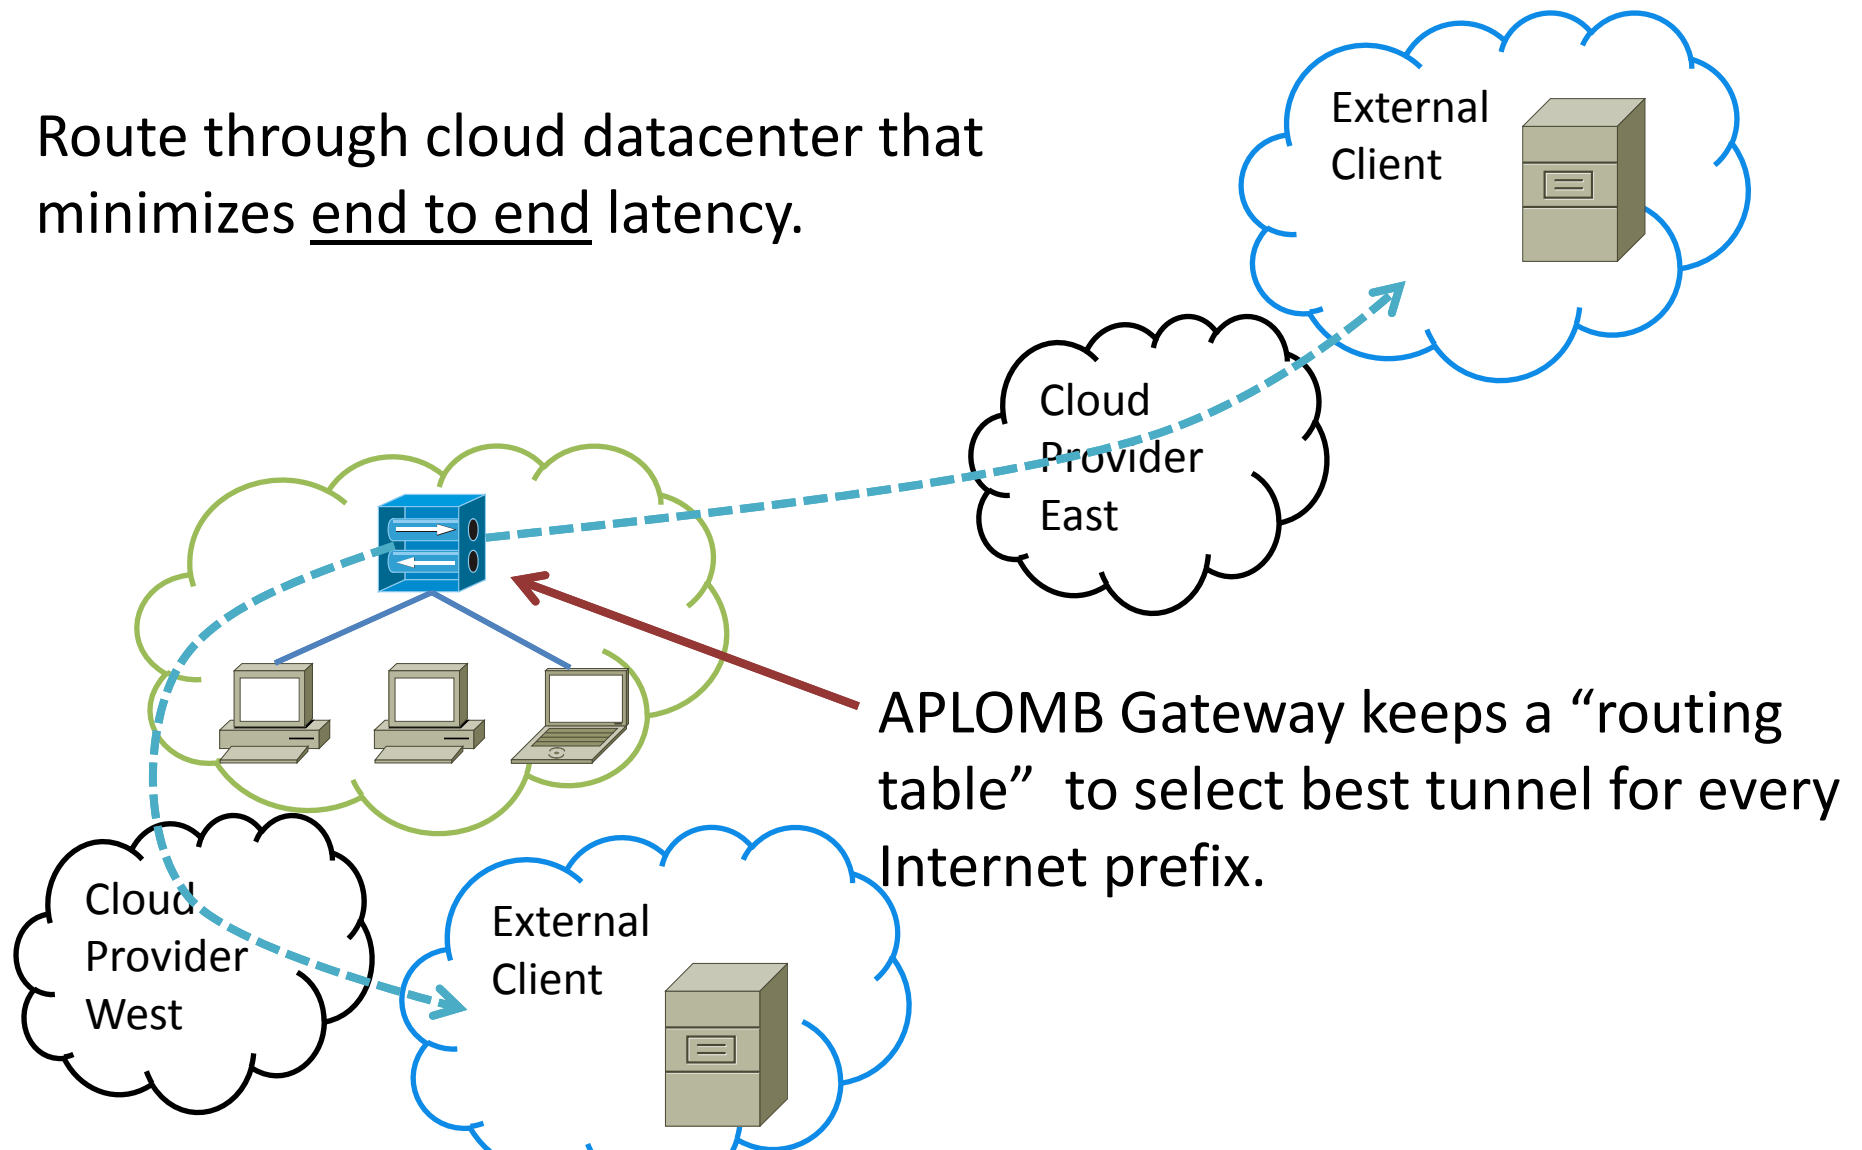
Minimizing latency?



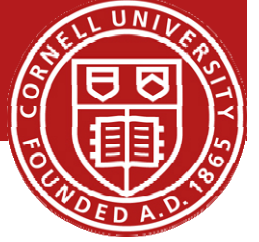
Choosing a Datacenter



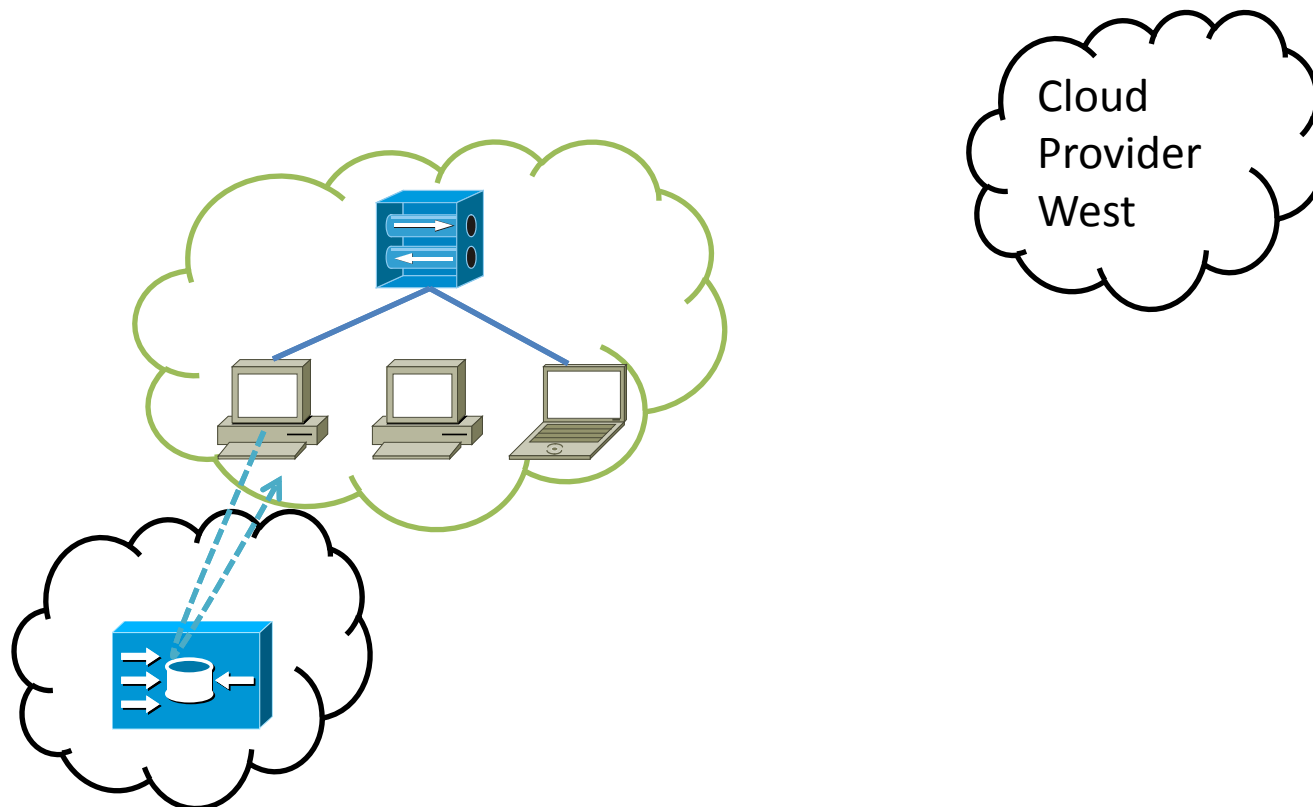
Route through cloud datacenter that minimizes end to end latency.



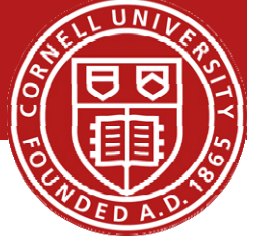
Caches and “Terminal Services”



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



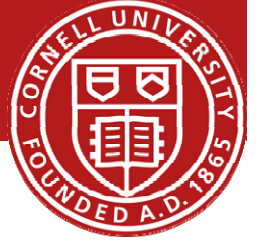
APLOMB



“Appliance for Outsourcing Middleboxes”

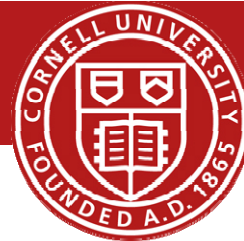
- Place middleboxes in the cloud.
- Use APLOMB devices and DNS to redirect traffic to and from the cloud.
- That’s it.

Can we outsource all middleboxes?

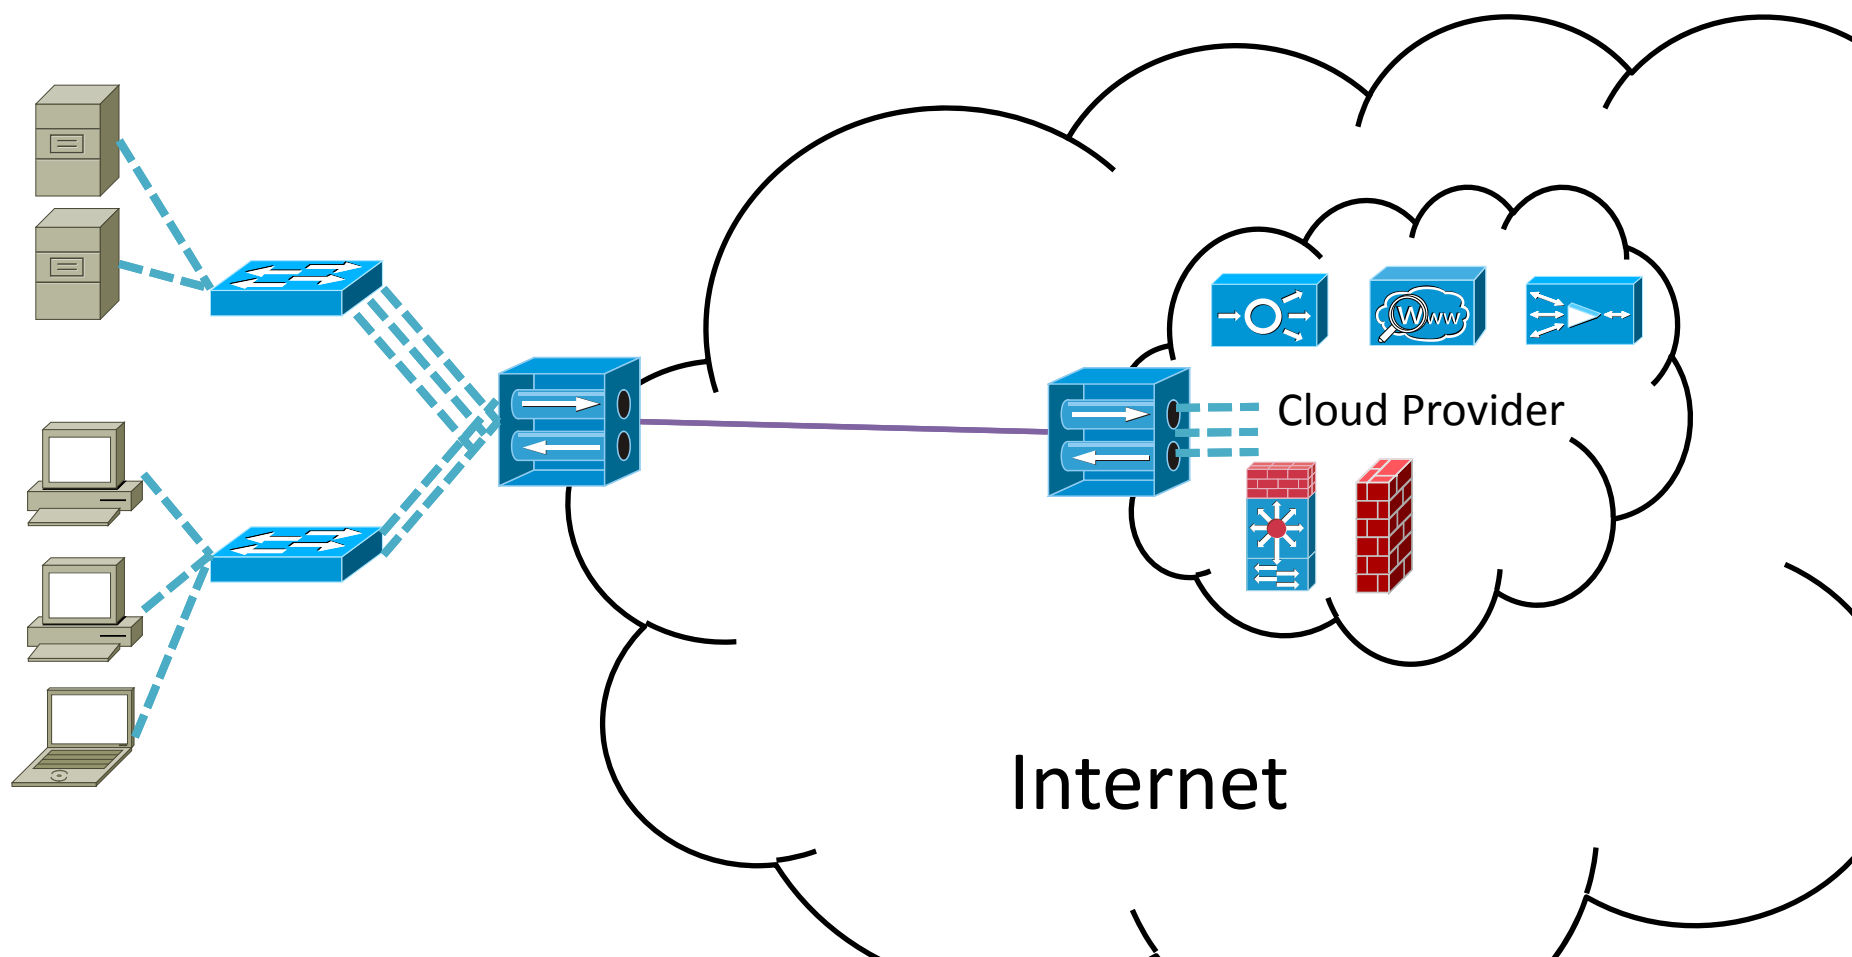


Firewalls	✓
IDSes	✓
Load Balancers	✓
VPNs	✓
Proxy/Caches	X Bandwidth?
WAN Optimizers	X Compression?

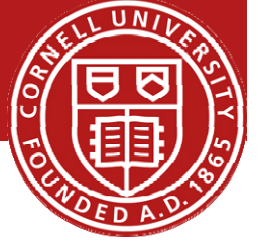
APLOMB+ for Compression



Add generic compression to APLOMB gateway to reduce bandwidth consumption.

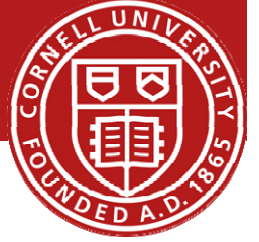


Can we outsource all middleboxes?

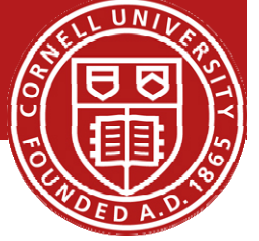


Firewalls	✓
IDSes	✓
Load Balancers	✓
VPNs	✓
Proxy/Caches	✗ Bandwidth?
WAN Optimizers	✗ Compression?

Does it work?

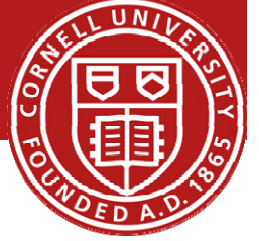


Deployment



- 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

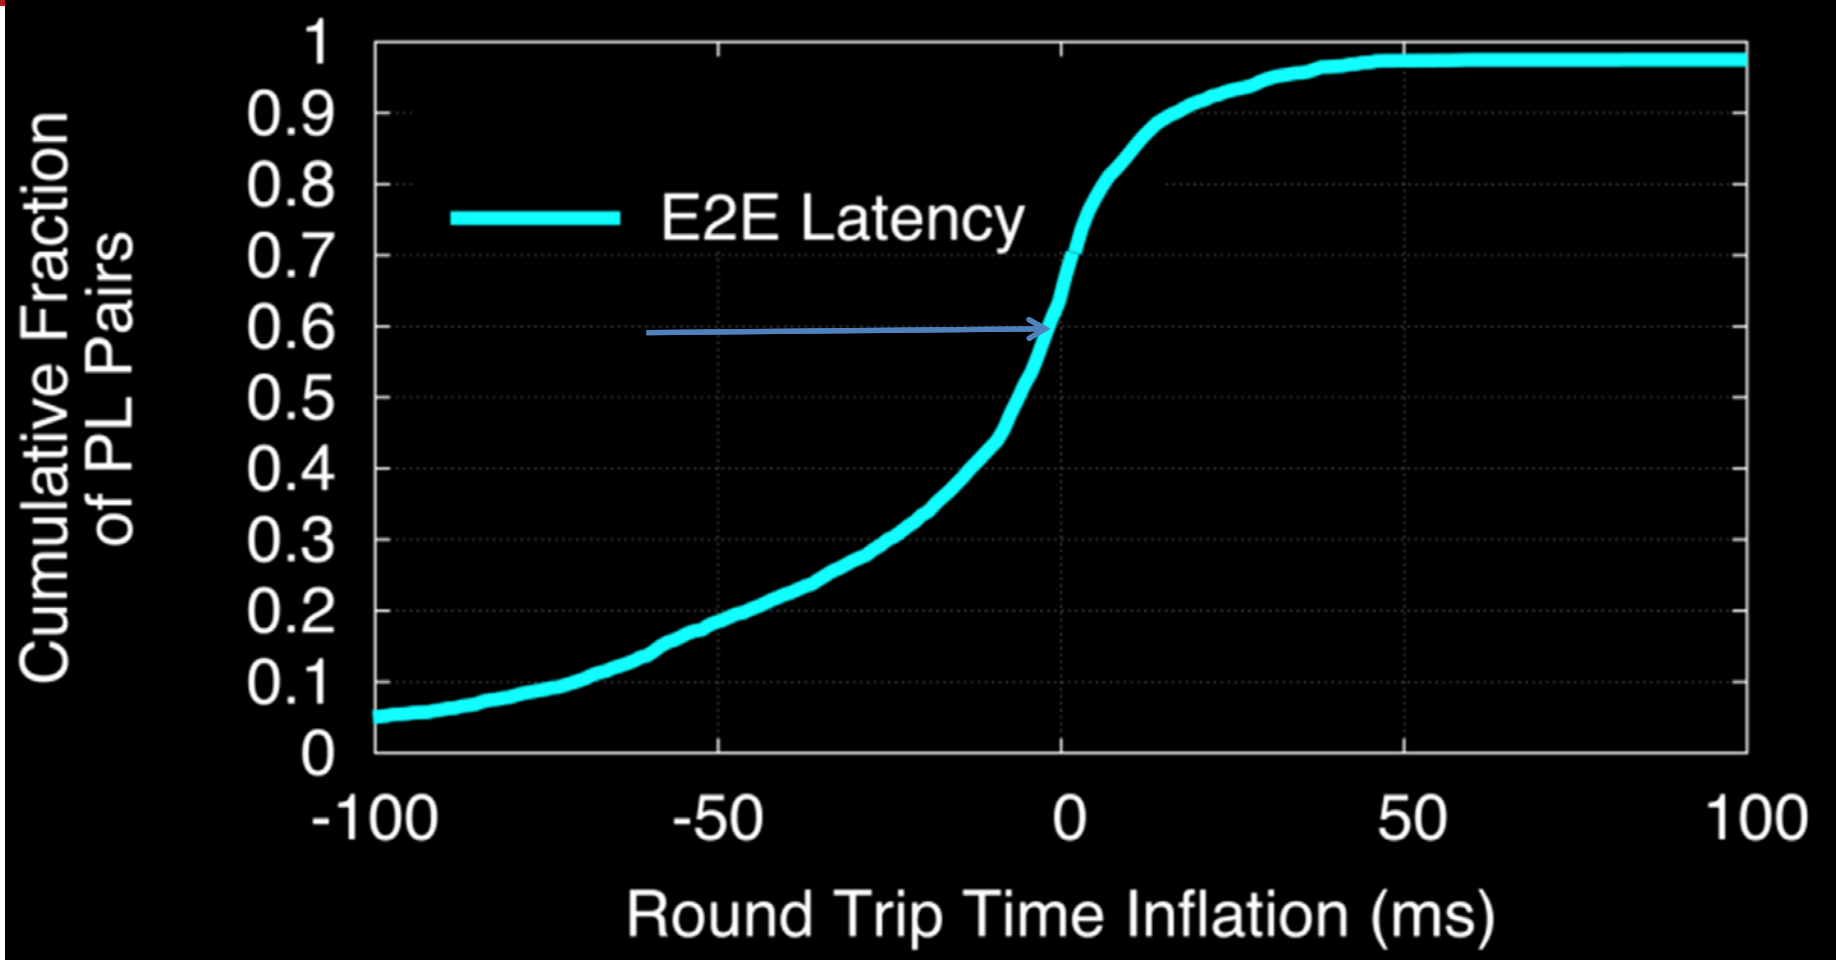
- Network latency

Case Study of a Large Enterprise

- Impact in a real usage scenario

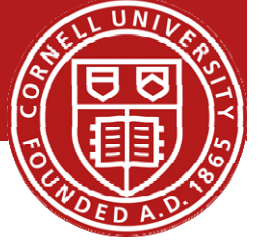
Does APLOMB inflate latency?





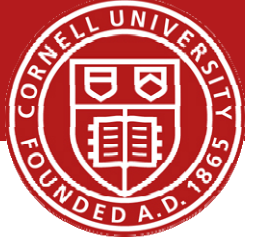
For PlanetLab nodes, 60% of pairs' latency improves with redirection through EC2.

Latency at a Large Enterprise

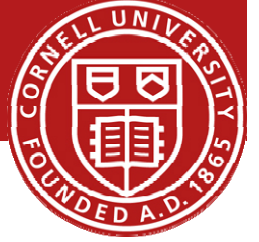


Measured redirection latency between enterprise sites.

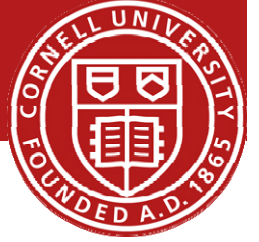
- Median latency inflation: 1.13 ms
- Sites experiencing inflation were primarily in areas where EC2 does not have a wide footprint.



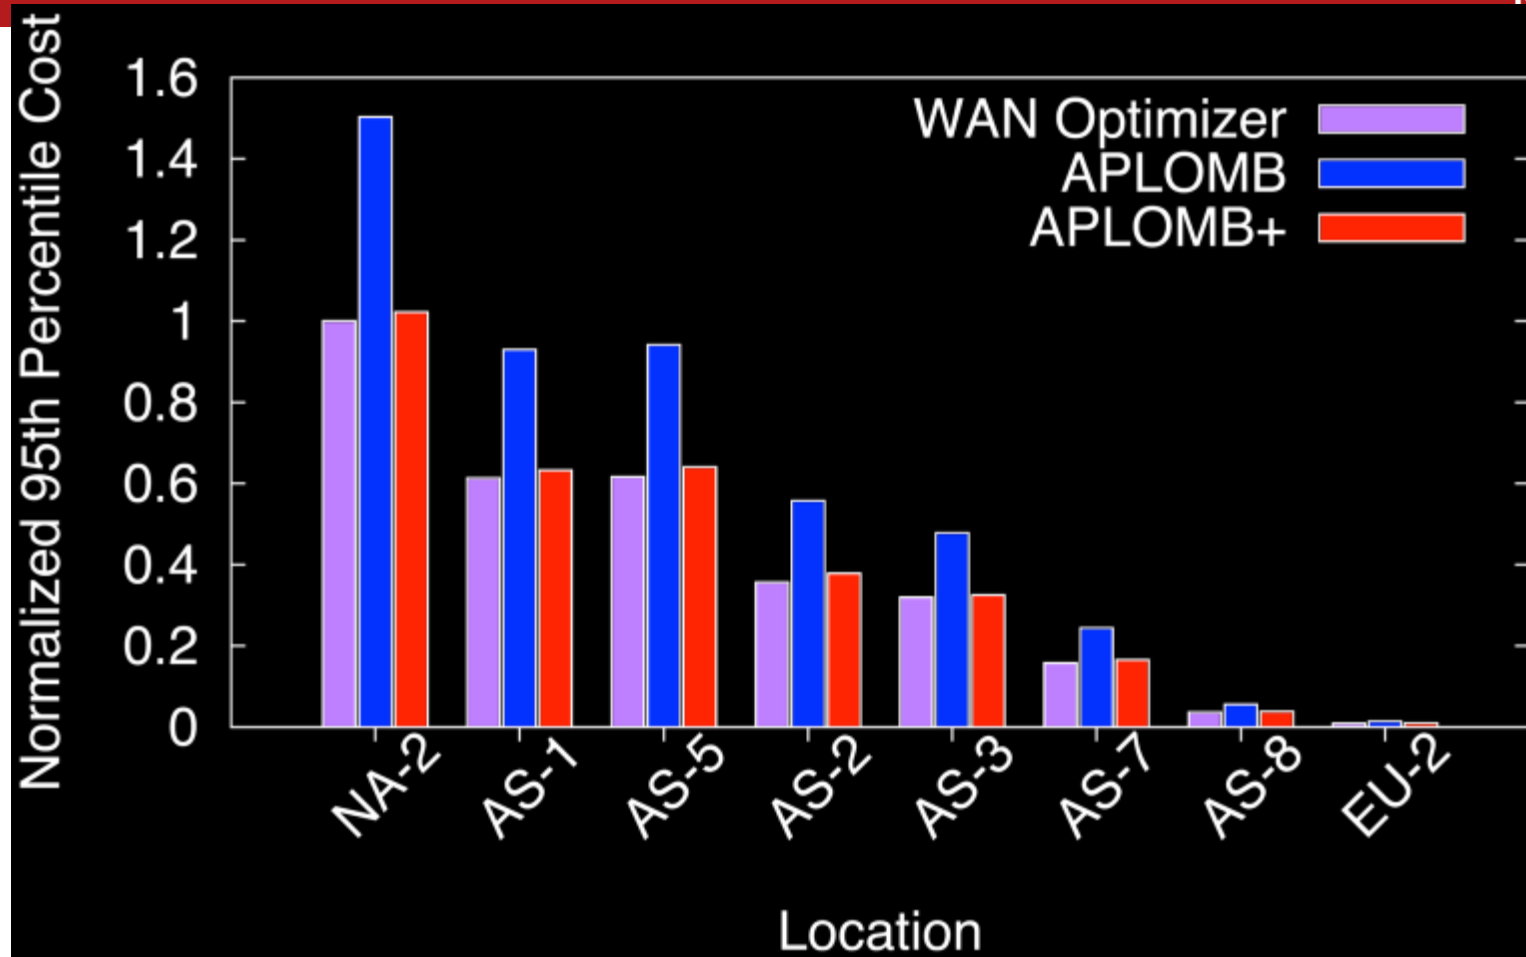
How does APLOMB impact other quality metrics, like bandwidth and jitter?



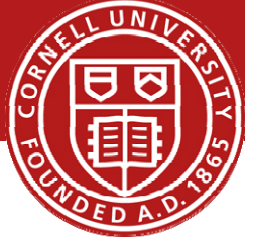
- **Bandwidth:** download times with BitTorrent increased on average 2.3%
- **Jitter:** consistently within industry standard bounds of 30ms



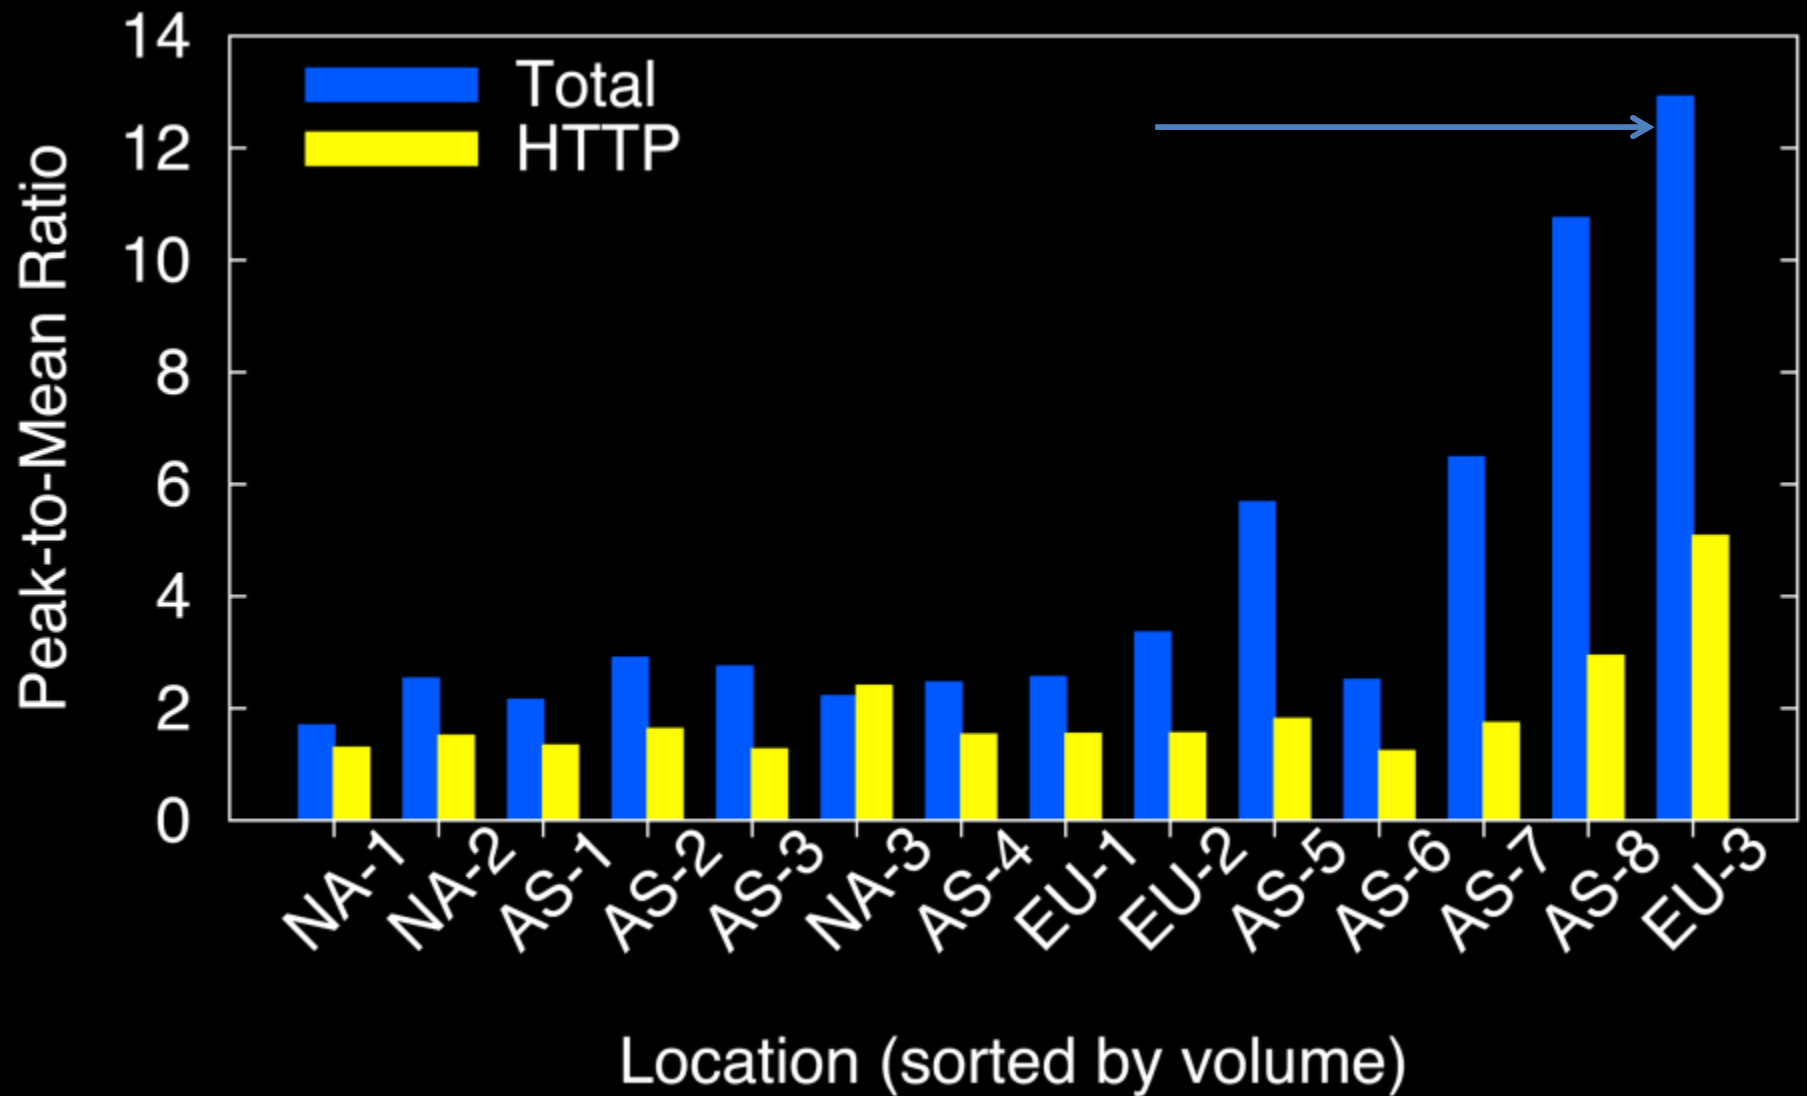
Does APLOMB negate the benefits of bandwidth-saving devices?



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

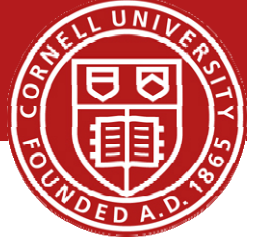


Does “elastic scaling” at the cloud
provide real benefits?



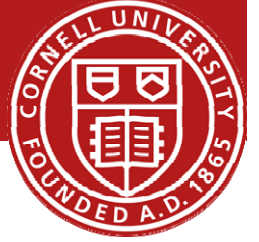
Some sites generate as much as 13x traffic more than average at peak hours.

Recap



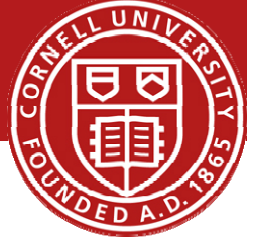
- Good application performance
 - Latency median inflation 1.1ms
 - Download times increased only 2.3%
- Generic redundancy elimination saves bandwidth costs
- Strong benefits from elasticity

Conclusion and Discussion



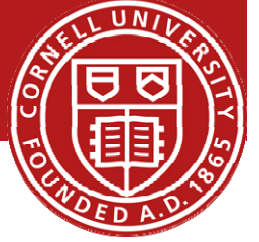
Moving middleboxes to the cloud seems to be **practical** and **feasible** solution to the complexity of enterprise networks.

Conclusion and Discussion



- Did the soln make the problem simpler?
 - How to measure simplicity/complexity?
- Does the soln also make security problems someone else's problems.
 - Do we trust the cloud provider?
- Privacy concerns?
 - Do we trust the cloud provider
- Monetary cost: Is APLOMB cheaper or more expensive?
- Precedence
 - Ariaka
 - Total uptime
- Middleboxes not at the edge of your network
 - APLOMB cannot outsource these middleboxes

Before Next time



- Project Interim report
 - Due **Today**, Monday, November 24.
 - And meet with groups, TA, and professor
- Fractus Upgrade: Should be back online
- ***Required review and reading for Monday, December 1***
 - IOFlow: a software-defined storage architecture, E. Thereska, H. Ballani, G. O'Shea, T. Karagiannis, A. Rowstron, T. Talpey, R. Black, T. Zhu. ACM Symposium on Operating Systems Principles (SOSP), October 2013, pages 182-196.
 - <http://dl.acm.org/citation.cfm?doid=2517349.2522723>
- Check piazza: <http://piazza.com/cornell/fall2014/cs5413>
- Check website for updated schedule