

Design and Implementation of a Consolidated Middlebox Architecture

Vyas Sekar

Sylvia Ratnasamy

Michael Reiter

Norbert Egi
Guangyu Shi



THE UNIVERSITY
of NORTH CAROLINA
at CHAPEL HILL




Need for Network Evolution

New applications



Evolving threats

Percentage of Methods Used to Exfiltrate Data



Policy constraints



New devices




Network Evolution today: Middleboxes!

<i>Type of appliance</i>	<i>Number</i>
Firewalls	166
NIDS	127
Media gateways	110
Load balancers	67
Proxies	66
VPN gateways	45
WAN Optimizers	44
Voice gateways	11
<i>Total Middleboxes</i>	<i>636</i>
<i>Total routers</i>	<i>~900</i>

Data from a large enterprise:
>80K users across tens of sites

Just network security
\$10 billion

Key pain points



Point solutions!



Increases capital expenses & sprawl
Increases operating expenses
Limits extensibility and flexibility


Outline

- Motivation
- *High-level idea: Consolidation*
- System design
- Implementation and Evaluation


Consolidation at Platform-Level

Today: Independent, specialized boxes


Proxy



Firewall




IDS/IPS



AppFilter




Decouple
Hardware and
Software



e.g., FlowStream
(UCL/Lancaster)


Consolidation reduces capital expenses and sprawl

Consolidation reduces CapEx




$$\text{Multiplexing benefit} = \frac{\text{Max_of_TotalUtilization}}{\text{Sum_of_MaxUtilizations}}$$

Consolidation Enables Extensibility




Contribution of reusable modules: 30 – 80 %

Consolidating Management



Consolidation enables flexible resource management

Today: All processing at logical "ingress"

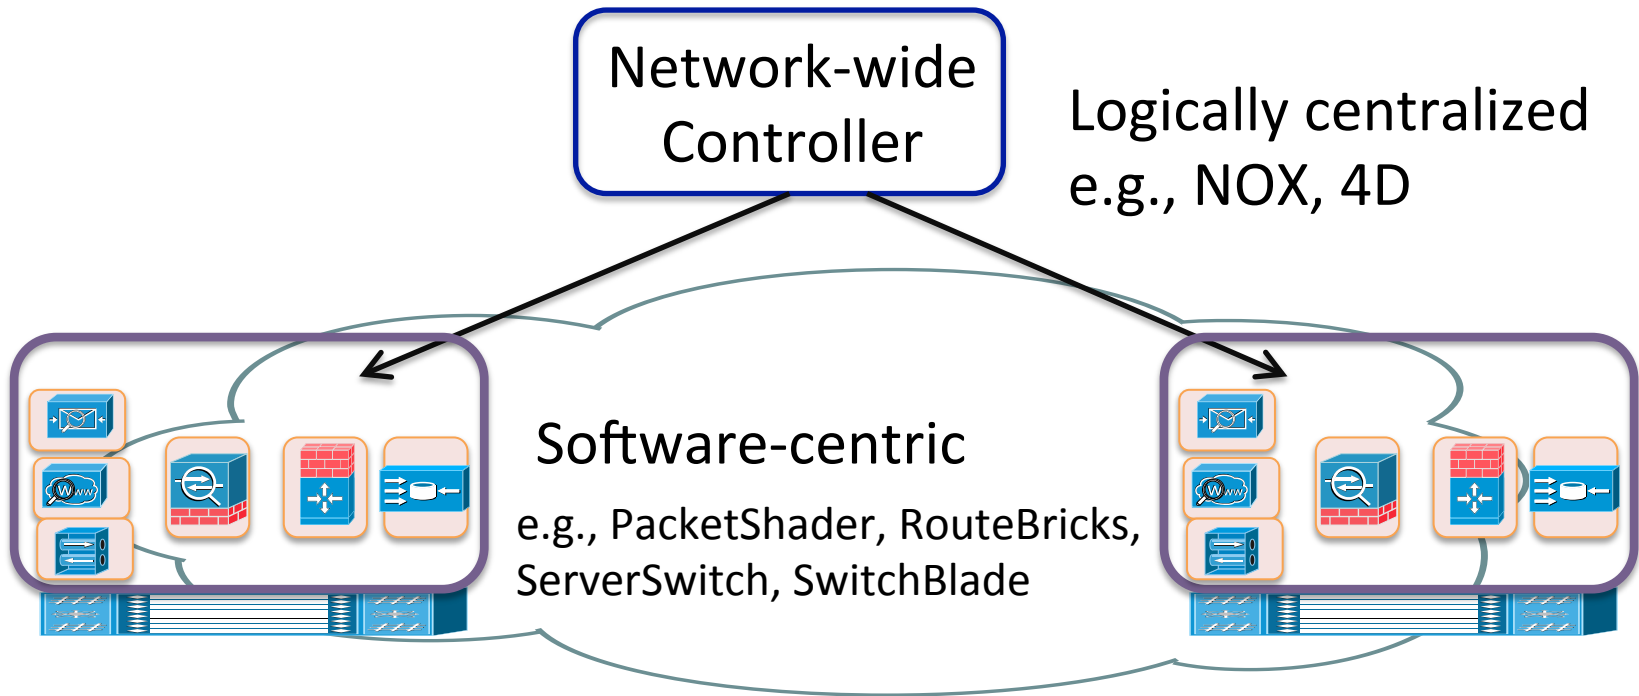


Distribution reduces load imbalance

Outline

- Motivation
- High-level idea: Consolidation
- *CoMb: System design*
- Implementation and Evaluation

CoMb System Overview




Existing work: simple, homogeneous routing-like workload

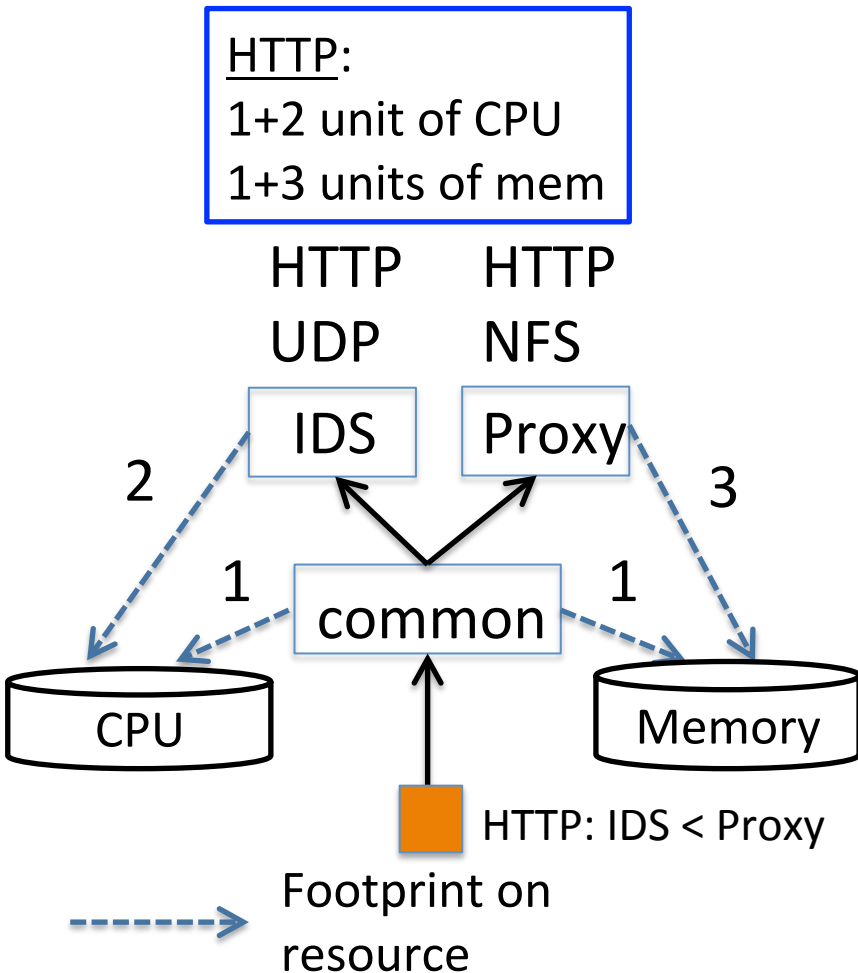
Middleboxes: complex, heterogeneous, new opportunities

CoMb Management Layer

Goal: Balance load across network
Exploit multiplexing, reuse, distribution




Capturing Policy and Reuse Efficiently



Need per-packet
 policy, reuse dependencies!

HyperApp: union of apps to run



Policy, dependency are implicit
 Needs small brute-force step

Network-wide Optimization

Minimize Maximum Load, Subject to

Processing coverage for each class of traffic

→ Fraction of processed traffic adds up to 1

*No explicit
Dependency
Policy*

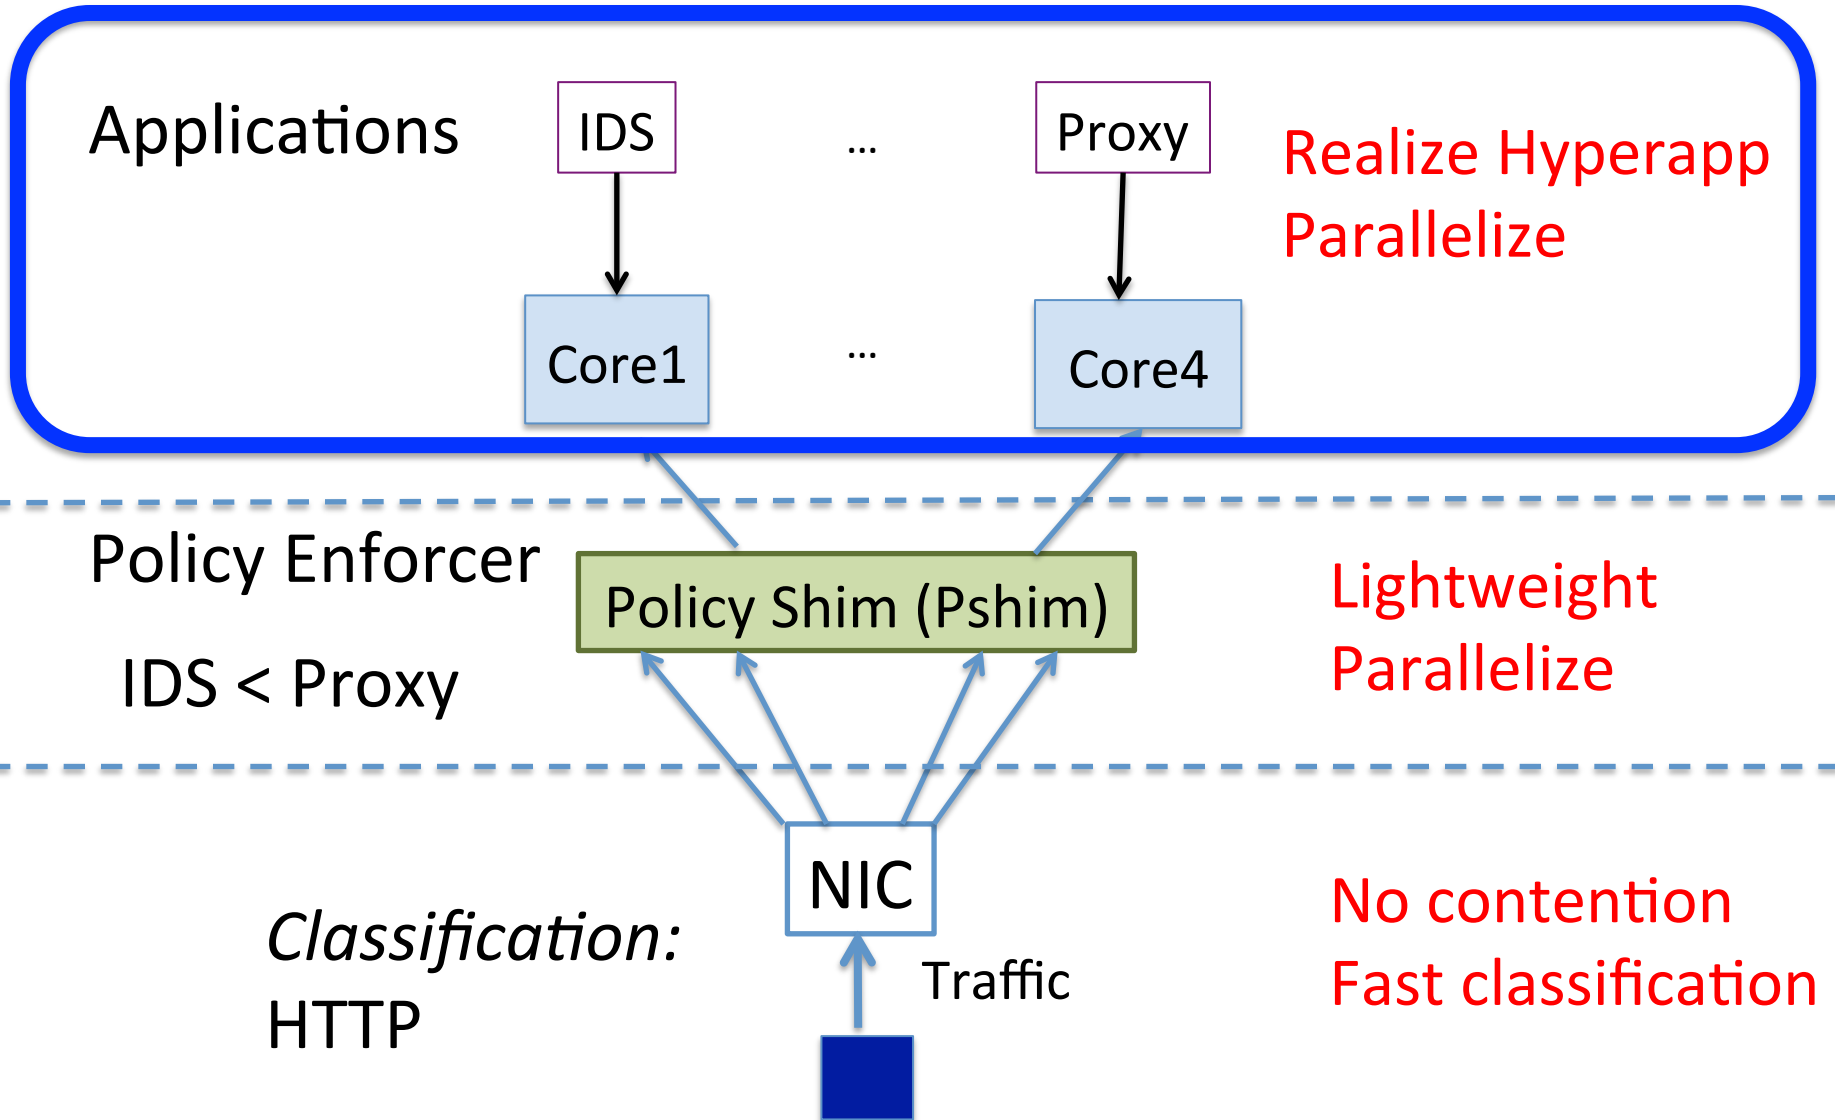
Load on each node

→ sum over HyperApp responsibilities per-path

A simple, tractable linear program

Very close (< 0.1%) to theoretical optimal

CoMb Platform

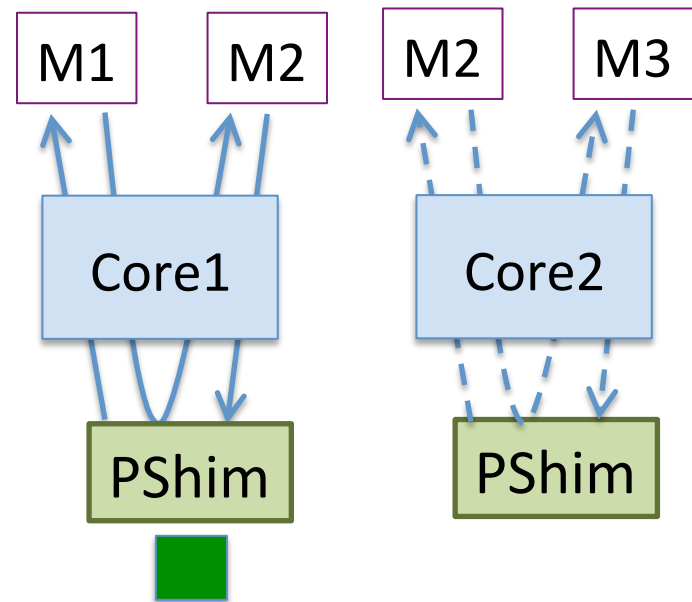



Parallelizing Application Instances

App
Per core

HyperApp1: $M1 < M2$
HyperApp2: $M2 < M3$

HyperApp
per core

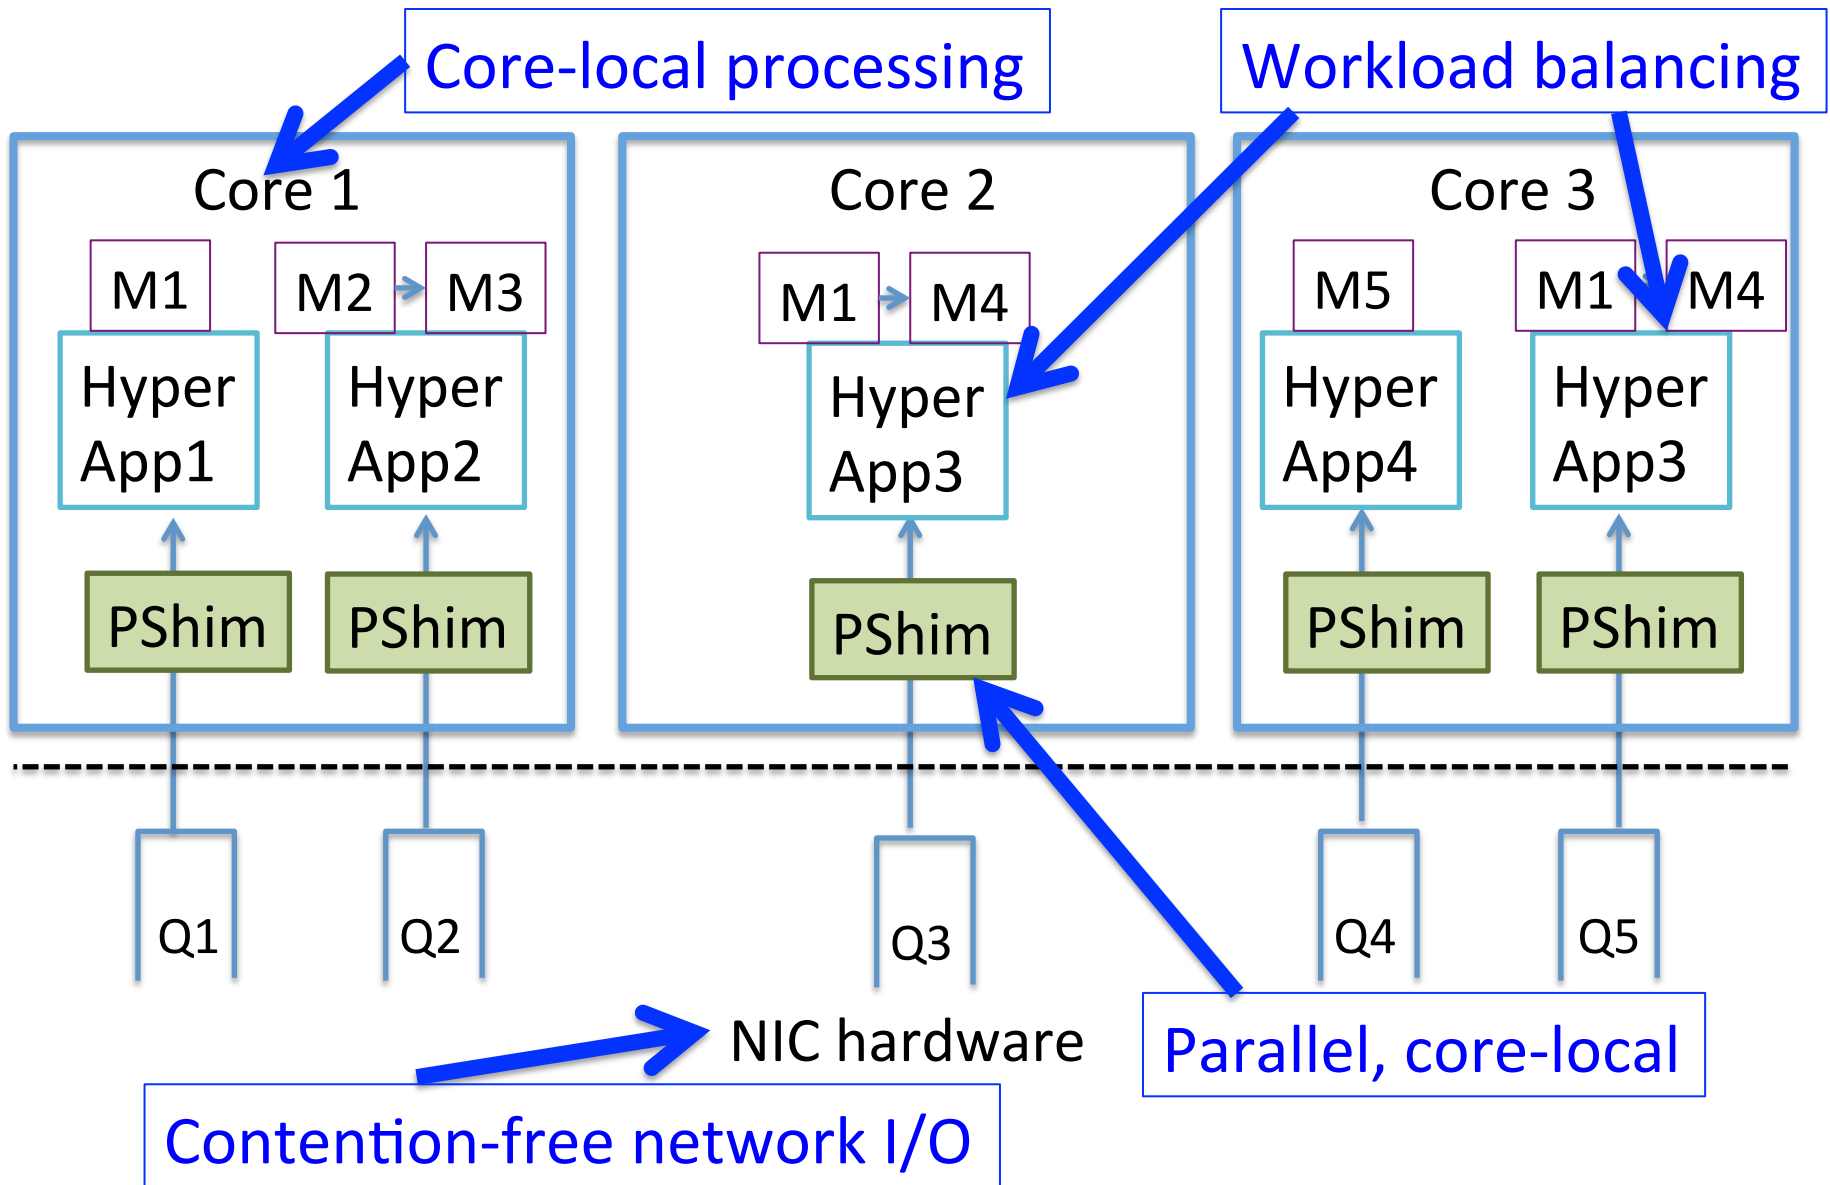


- Inter-core communication
- More work for PShim
- + No in-core context switch

- + Keeps structures core-local
- + Better for reuse
- But incurs context-switch
- Need replicas

HyperApp-per-core is better or comparable

CoMb Platform Design



Outline

- Motivation
- High-level idea: Consolidation
- System design: Making Consolidation Practical
- *Implementation and Evaluation*

CoMb Implementation

Network-wide Management

*using
CPLEX*




Policy Shim

Kernel mode Click

Extensible apps

Standalone apps




Protocol

Session

*Ported logic
From
Bro → Click*

*Memory mapped
Or
Virtual interfaces*

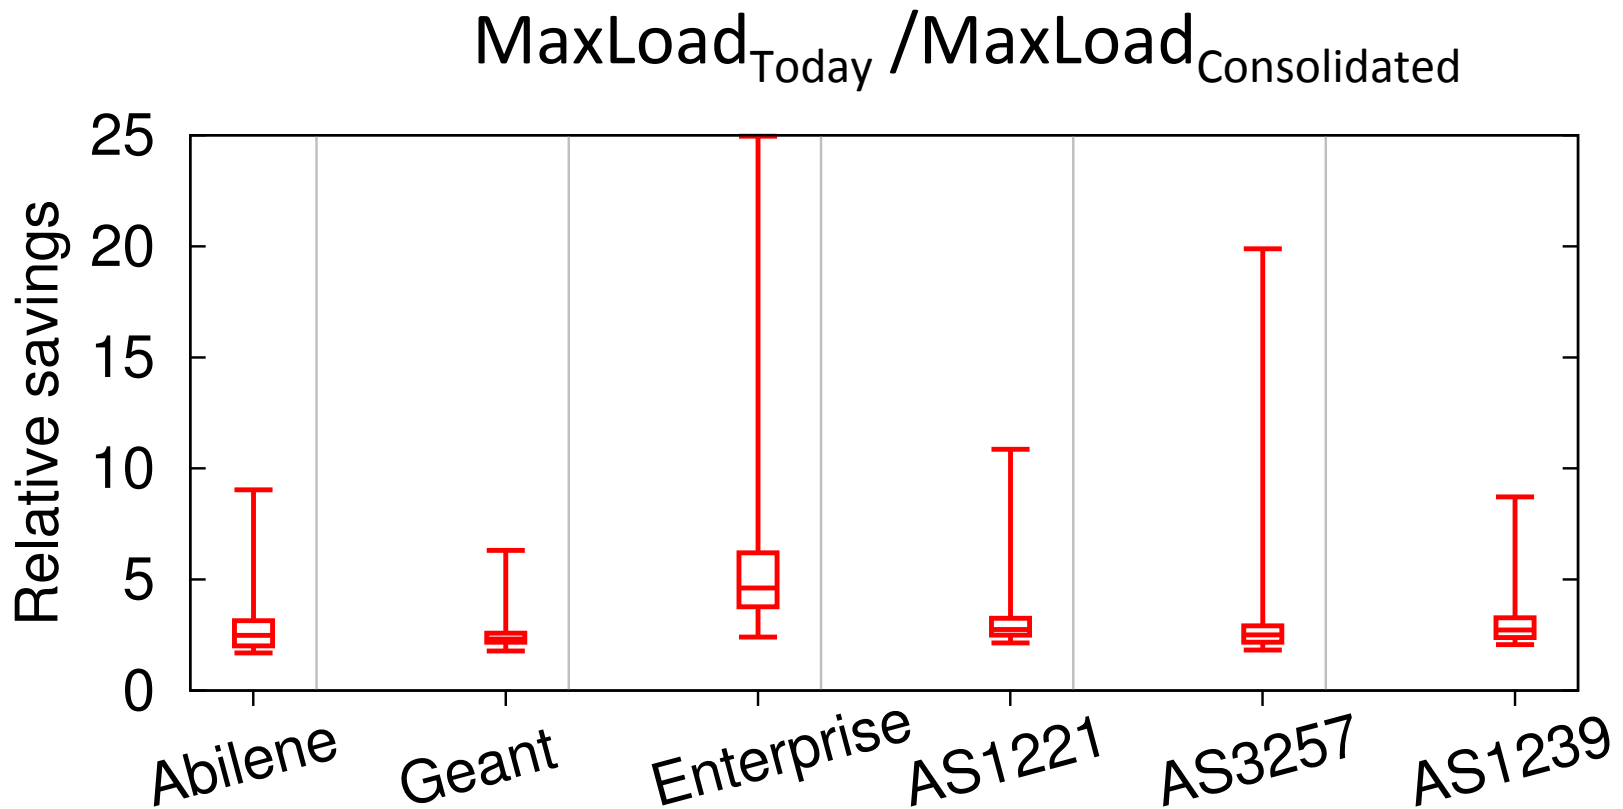


8-core Intel Xeon with Intel 82599 NIC

Consolidation is Practical

- Low overhead for existing applications
- Controller takes $< 1.6s$ for 52-node topology
- 5x better than VM-based consolidation

Benefits: Reduction in Maximum Load



Consolidation reduces maximum load by 2.5-25X

Consolidation reduces provisioning cost 1.8-2.5X

Discussion

- Isolation
 - Current: rely on process-level isolation
 - Leverage “user-space” networking
 - Get reuse-despite-isolation?
- Changes vendor business models
 - Already happening (e.g., “virtual appliances”)
 - Benefits imply someone will do it!
 - May already have extensible stacks

Conclusions

- Most network evolution today occurs via middleboxes
- Today: Narrow, point solutions
 - High CapEx, OpEx, and device sprawl
 - Inflexible, difficult to extend
- Our proposal: Consolidated architecture
 - Extensible, general-purpose
 - Reduces CapEx, OpEx, and device sprawl
- More opportunities
 - Isolation
 - APIs (H/W—Apps, Management—Apps, App Stack)