15-849 Datacenter Computing

Dimitrios Skarlatos

Fall 2021

Agenda

- Introductions
- Logistics
- Course Overview

Introductions

Schedule Overview

Logistics

Administrivia

Instructor

- Dimitrios Skarlatos
  dskarlat@cs.cmu.edu

Course Website

- https://www.cs.cmu.edu/~15849
- Schedule
- Syllabus
- Assignments

Office Hours

- Tuesdays 11am

Piazza

- Announcements
- Online discussion

HotCRP (yes!)

- Paper reviews
- Act like PC members*

Canvas

- Project proposal and report

*without rejections

Prerequisites

System's related doctoral students

- 15-213 Introduction to Computer Systems

Recommended

- 15-740 Computer Architecture
- 15-410 Operating System Design and Implementation
- 15-316 Software Foundations of Security and Privacy

Grading

Lab Assignment: 15%

Seminar: 35%

- Paper reviews (2 per week) - 20%
- Presentation lead - 10%
- Participation/Discussion (2 papers per course) – 5%

Research Project: 50%

- Proposal - 10%
- Midterm Checkpoint - 10%
- Final Presentation - 15%
- Final Report - 15%

Bonus

- Project gets accepted to a top tier conference you automatically get an A!

Lab Assignment

Will be announced last week of September

- Due mid-October (~about three weeks)

Teams of two

- Use piazza to find teammates

Lab Assignment

- Safecracker!

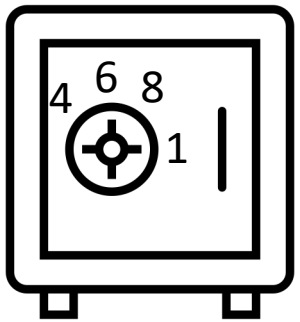

Lab Assignment

- Safecracker!

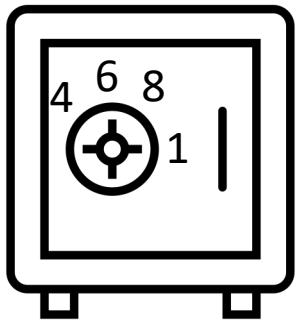

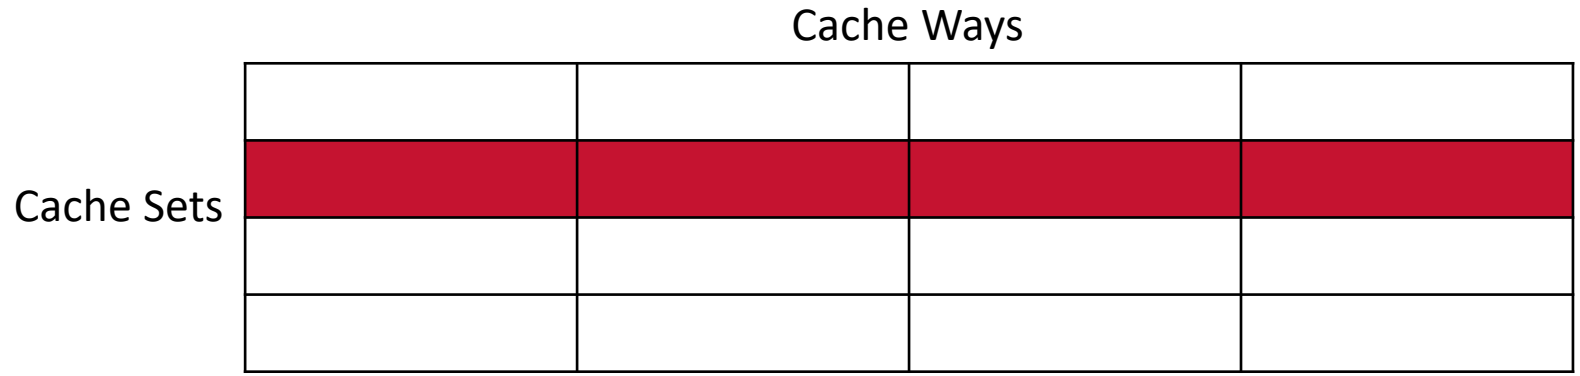

Victim

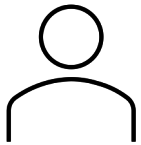

Safe combination in set X
open safe → access set X

t = time(setX_access)
if t > slow_threshold
  combination (X)
else
  check other set

Thief

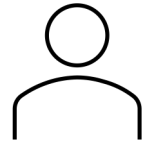

Lab Assignment

Communicate through microarchitectural contention!

- No sockets
- No IPC

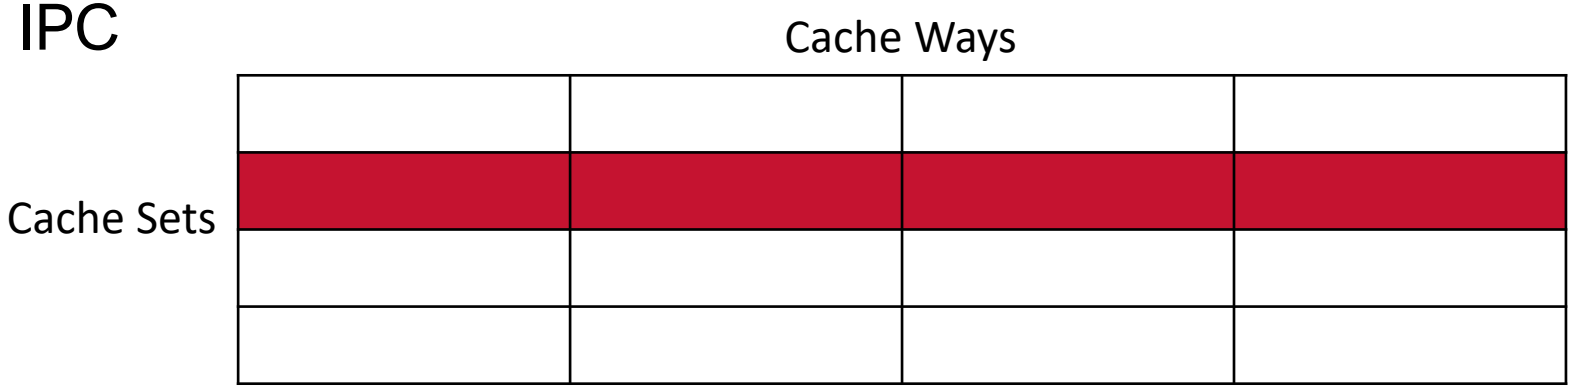

Sender

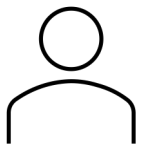

If Send 1
  fill set
else
  do nothing

t = time(set_access)
if t > slow_threshold
  receive (0)
else
  receive(1)

Receiver

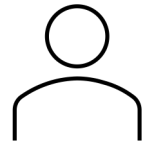

Paper Review Guidelines

Two papers per week

Due the midnight before each class

Submit review in HotCRP

Reviews will be made visible to everyone after the deadline

Reviews Format

A summary of the main idea
- Three sentences

Opportunities & challenges of a cross-layer design
- Three sentences

Limitations
- Two sentences

Potential Future Work
- Two sentences

Three Questions for discussion

Presentation Guidelines

About three presentations in the semester

Rank 6 papers ~2 from each month

- One topic you understand very well
- One topic you have no idea about
- One in between

Submit preferences on Google Form:

- Form link

Preliminary slides due 24 hours before class

- Can continue editing until the class

Presentation Format

Two presentations per course

- ~25-30 minutes for each presentation
- ~10-15 minutes for discussion

Slides should cover:

- Background and motivation
- Key technical contributions and design
- Strengths and weaknesses
- Opportunities & challenges for a cross-layer design
- Directions for future work
- Several questions for discussion (~5-10)

Project

Original research project with focus on cross-layer design

Teams of two or solo

Pre-proposal

- Schedule meeting with me in the next 2 weeks

Proposal

- 1 page summary
- Motivation, high-level design, expected evaluation

Checkpoint

- Schedule meeting in mid-October
- 5-minute presentation

Final presentation and report

- December 2 and 7, no finals!

Bonus

Project gets accepted to a top tier conference you automatically get an A!

Project Ideas

Bridge Operating Systems and Hardware

Performance Security Scalability

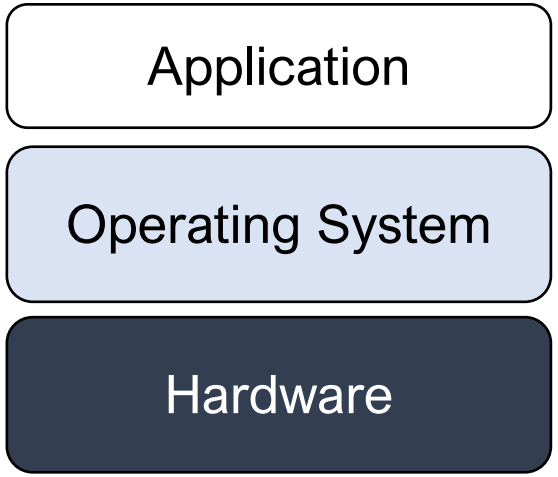

Virtualization

Serverless

FaaS Storage

Networking

Operating Systems

Containers Microservices

Processor Pipeline

Caches DRAM/NVM

Accelerators

Enclaves

Side-channels

Attacks & Defenses

Speculative Execution

System-level security

Academic Integrity & Collaboration

Discussion of research papers and ideas is strongly encouraged

Paper reviews must be done individually and on your own writing

Teams may use any online code and material with clear citations

Teams must provide their own proposal, presentations, code, & report

CMU's Academic Integrity Policy

Student Wellness

You may face a range of challenges that may affect your academic performance and daily life.

CMU offers Counseling and Psychological Services (CaPS)

The Computer Architecture Student Association (CASA) is an independent student-run organization with the express purpose of developing and fostering a positive and inviting student community within computer architecture

Action Items for Today

Check paper schedule

- https://www.cs.cmu.edu/~15849/schedule.html

Fill preference form

- https://forms.gle/JZ93UQvwtepL9KKm7

Course Overview

Datacenter Computing

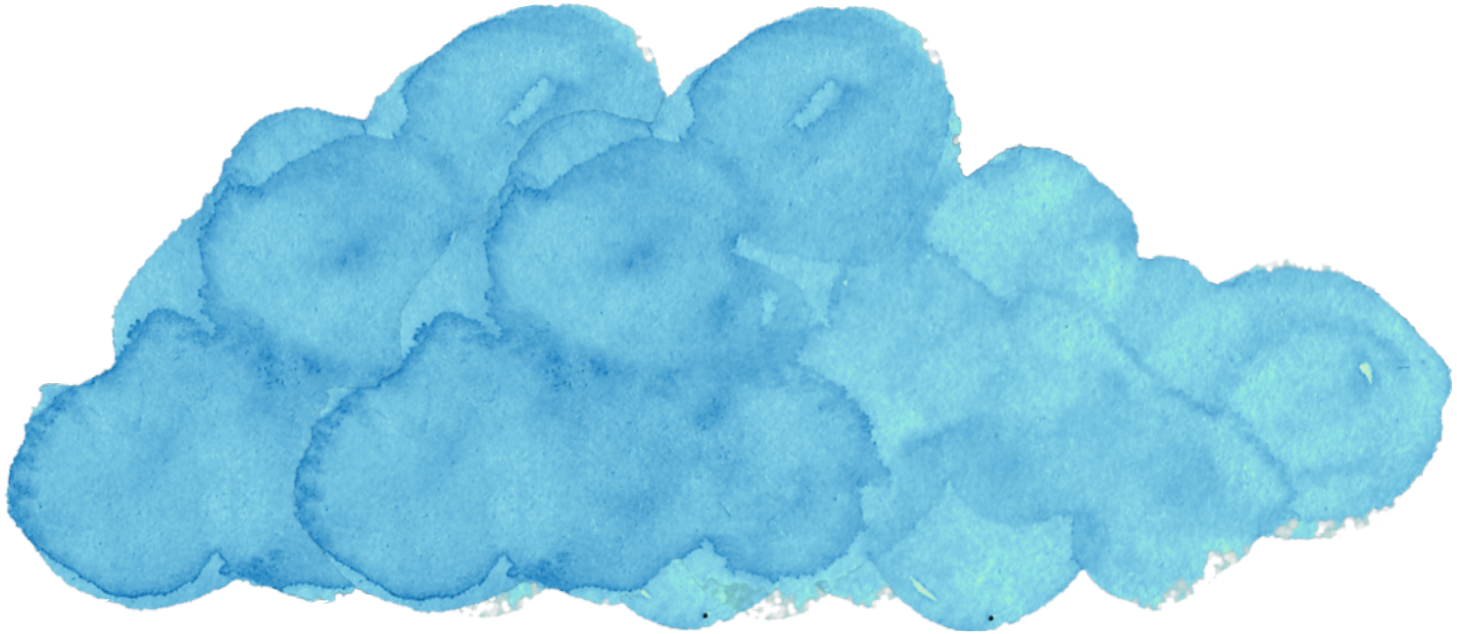

Datacenter Computing

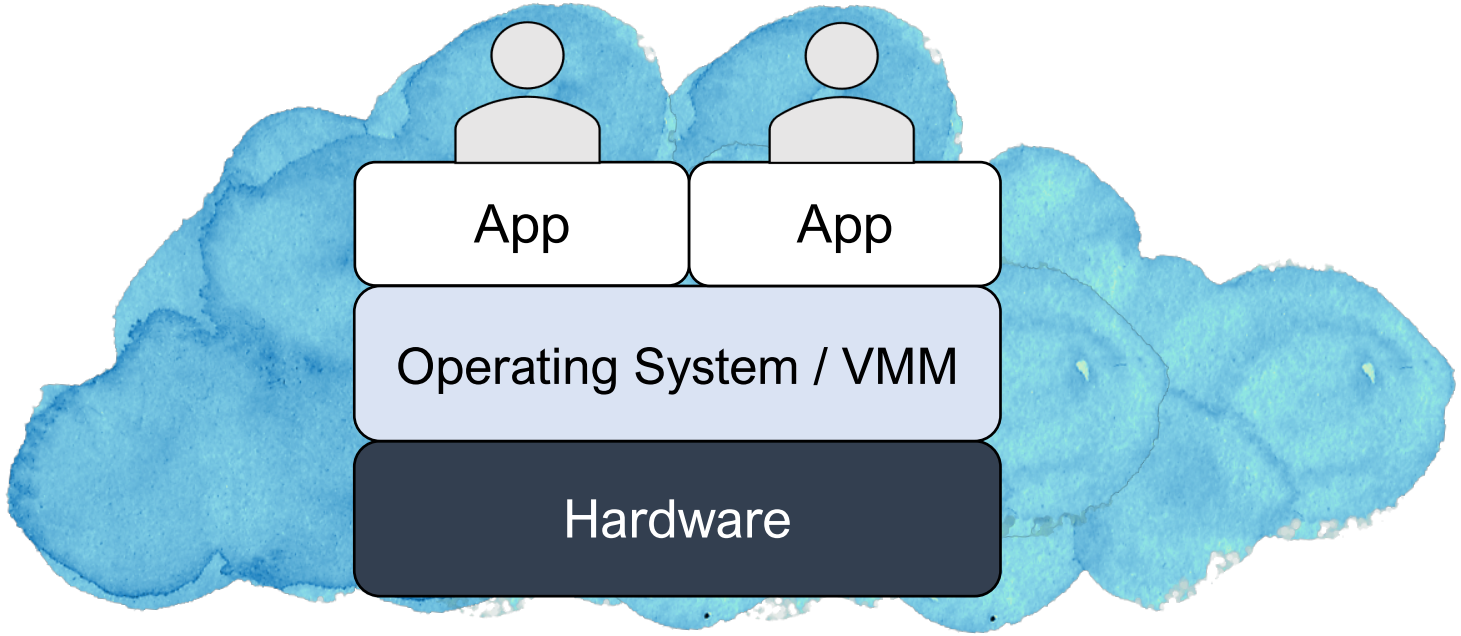

Radical Shift in Datacenter Computing

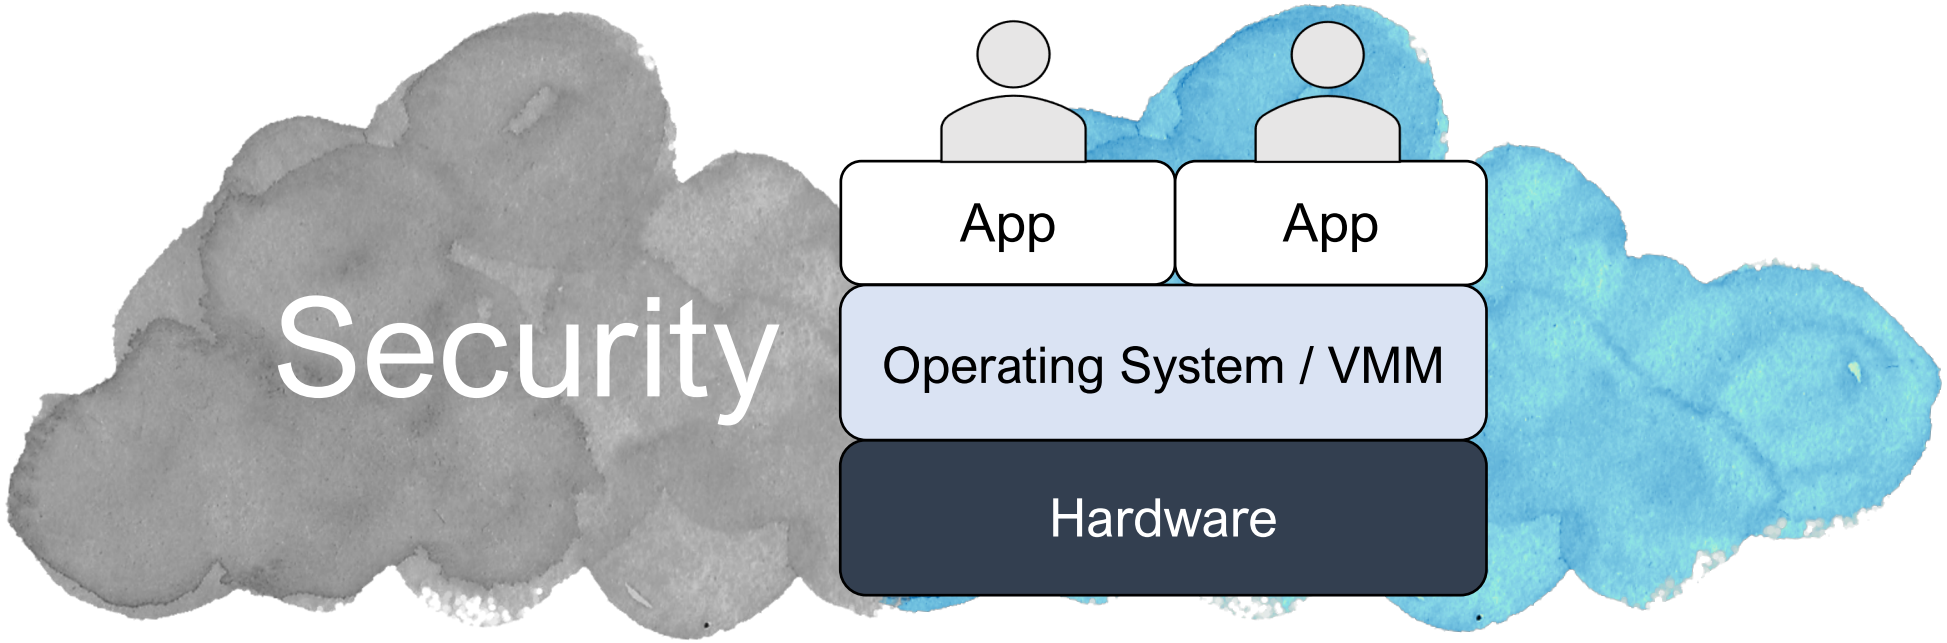

Radical Shift in Datacenter Computing

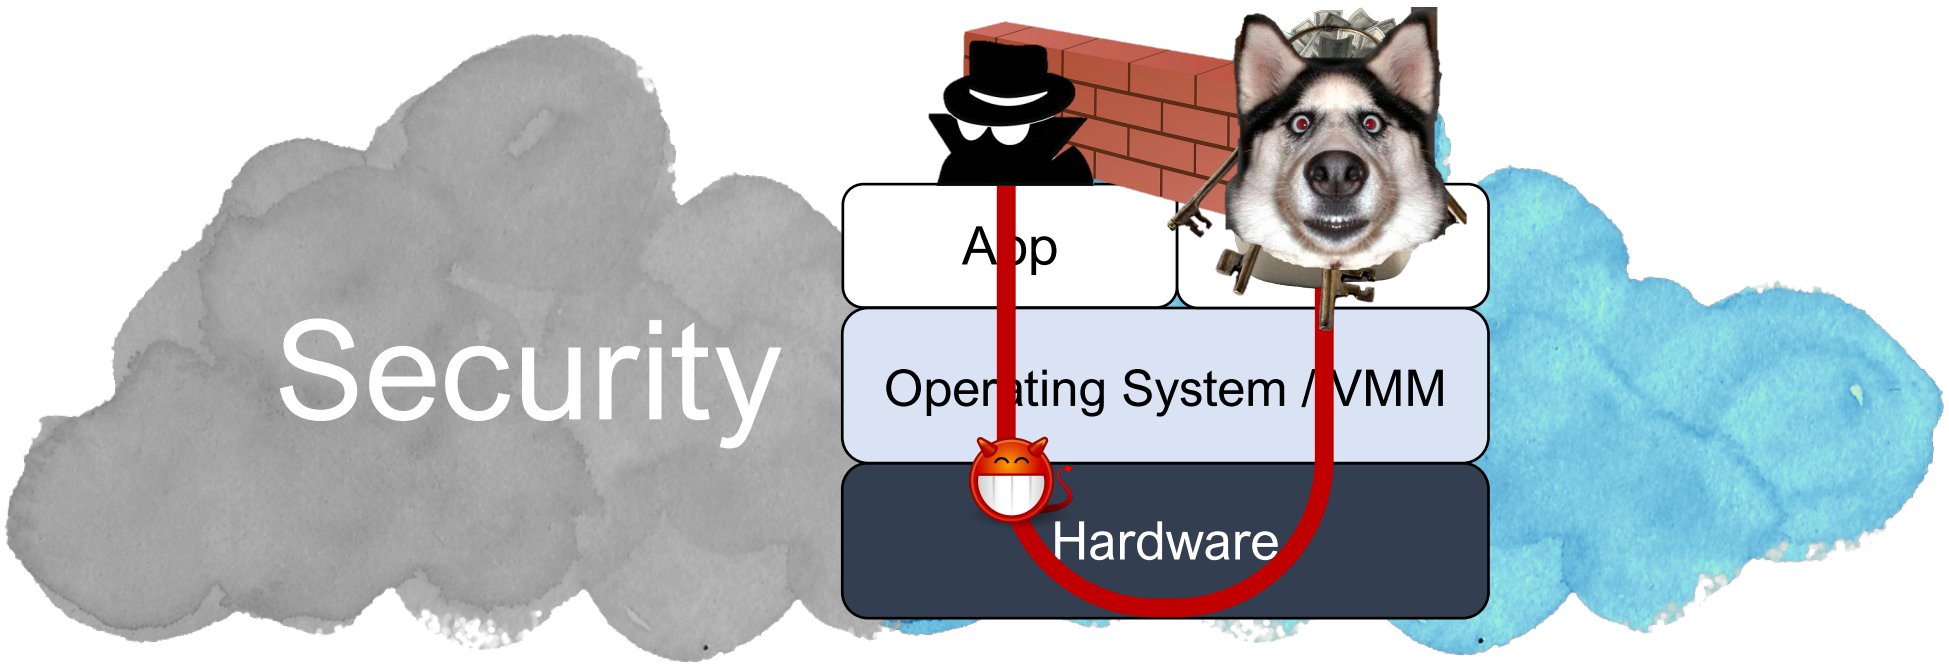

Security is a critical design requirement

- System attack vectors
- Side-channel attacks
- Bypass current system barriers

Radical Shift in Datacenter Computing

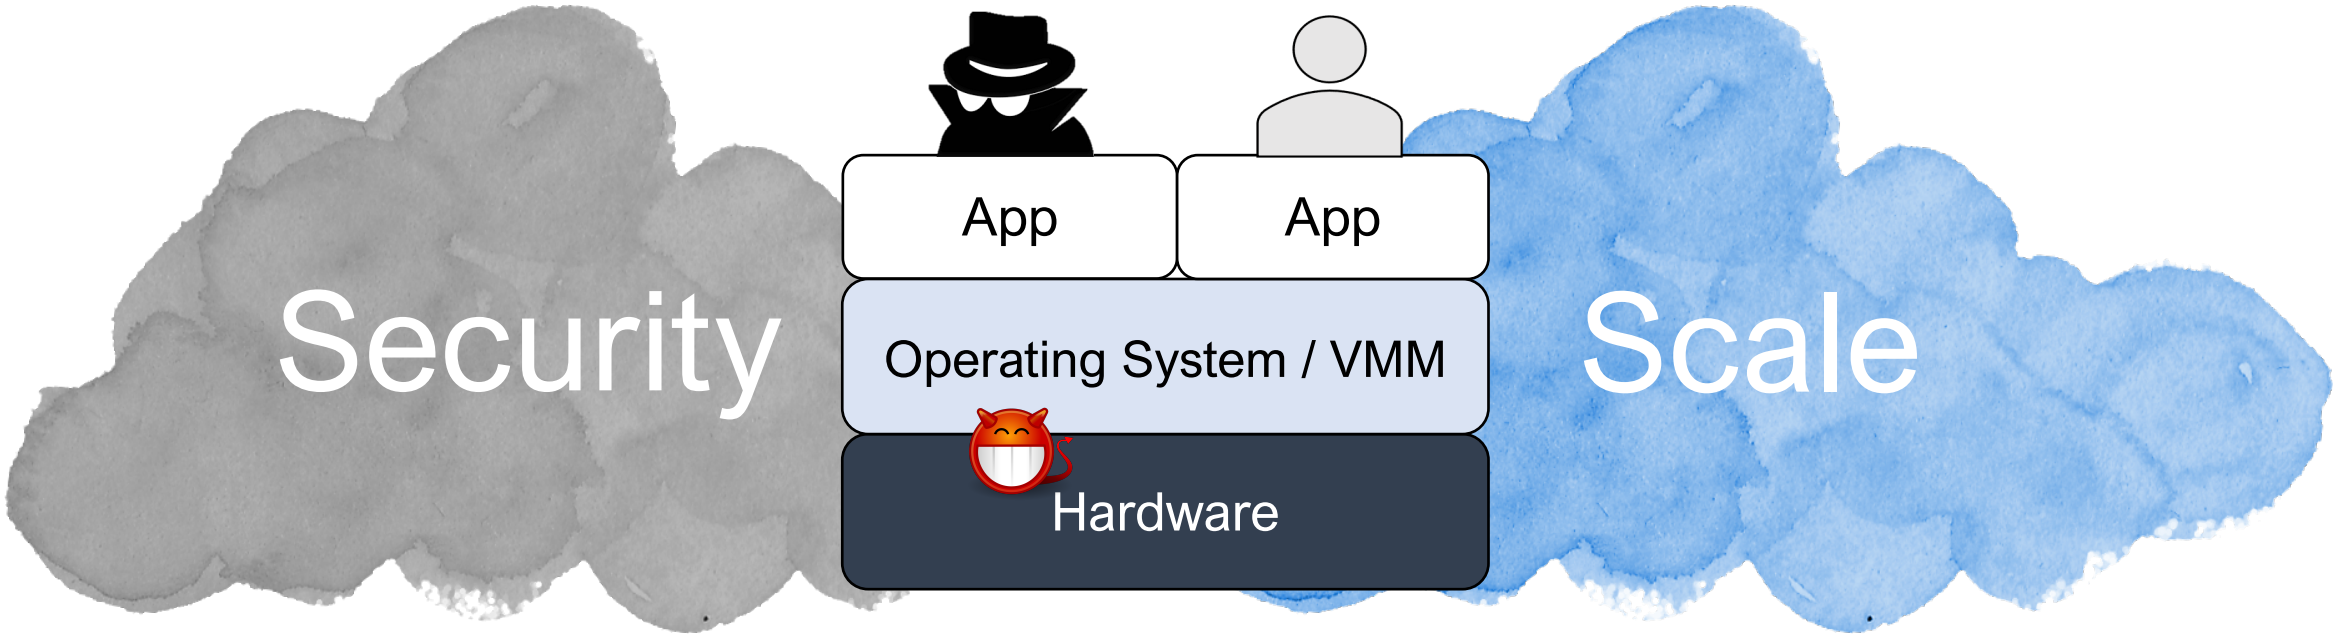

Security is a critical design requirement

- System attack vectors
- Side-channel attacks
- Bypass current system barriers

Unprecedented growth in data and users

Emerging computing paradigms

- Microservices
- Serverless (Function-as-a-Service)

Radical Shift in Datacenter Computing

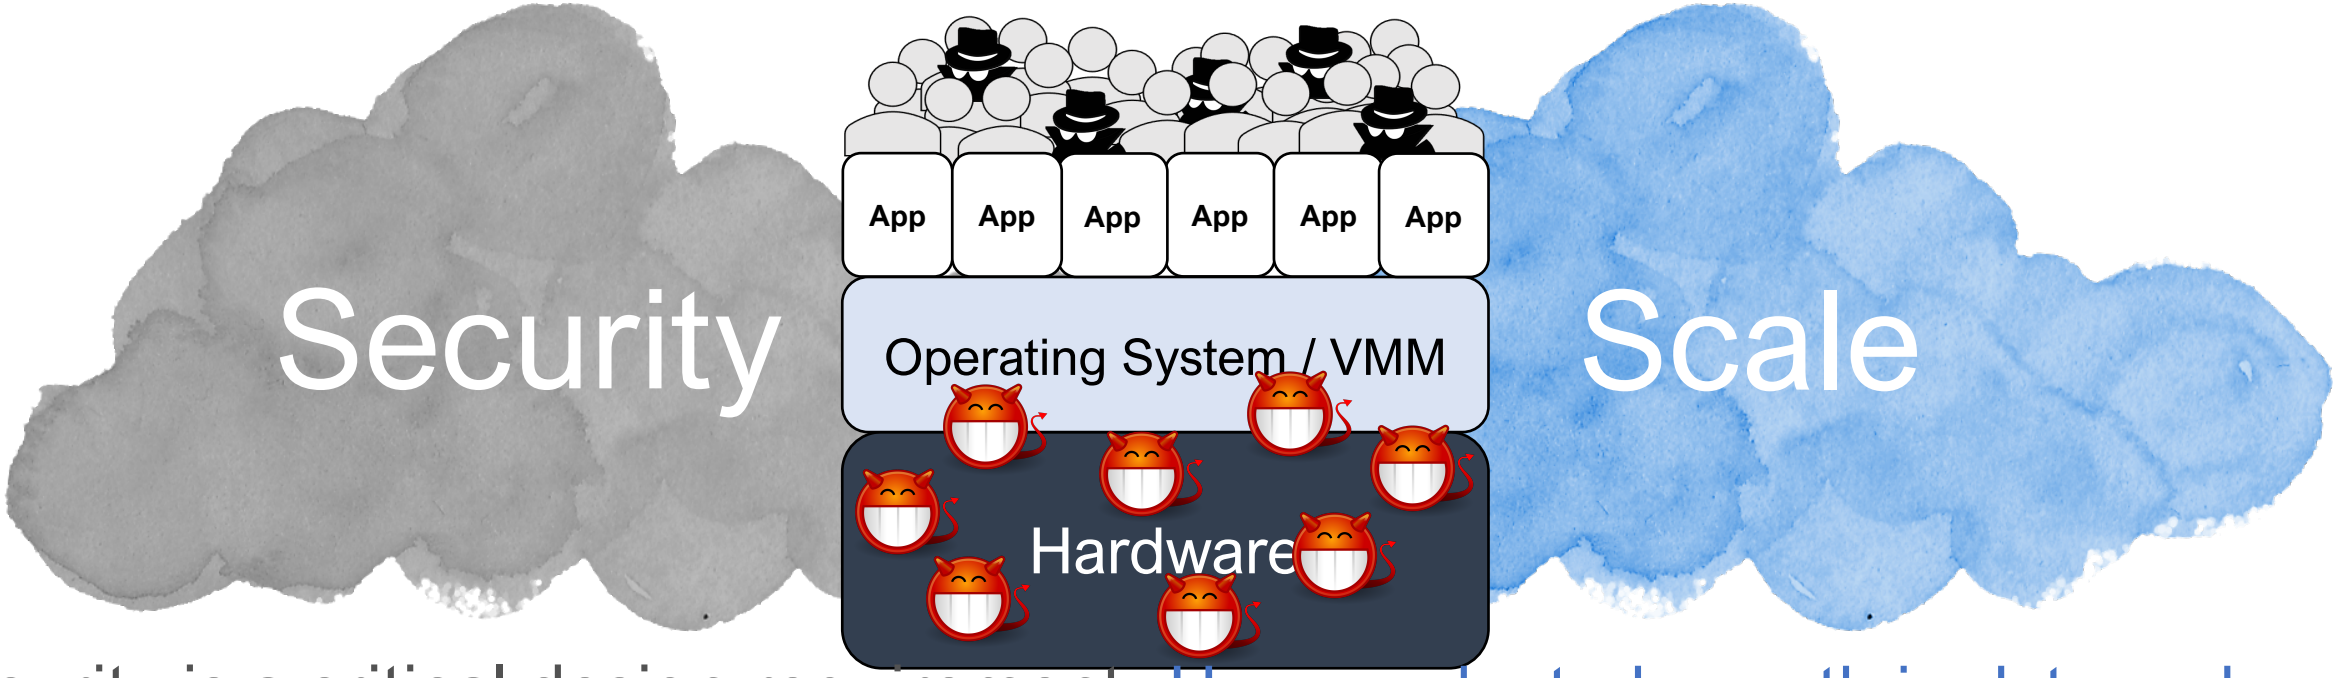

Security is a critical design requirement

- System attack vectors
- Side-channel attacks
- Bypass current system barriers

Unprecedented growth in data and users

Emerging computing paradigms

- Microservices
- Serverless (Function-as-a-Service)

Reactive Research

Reactive Research

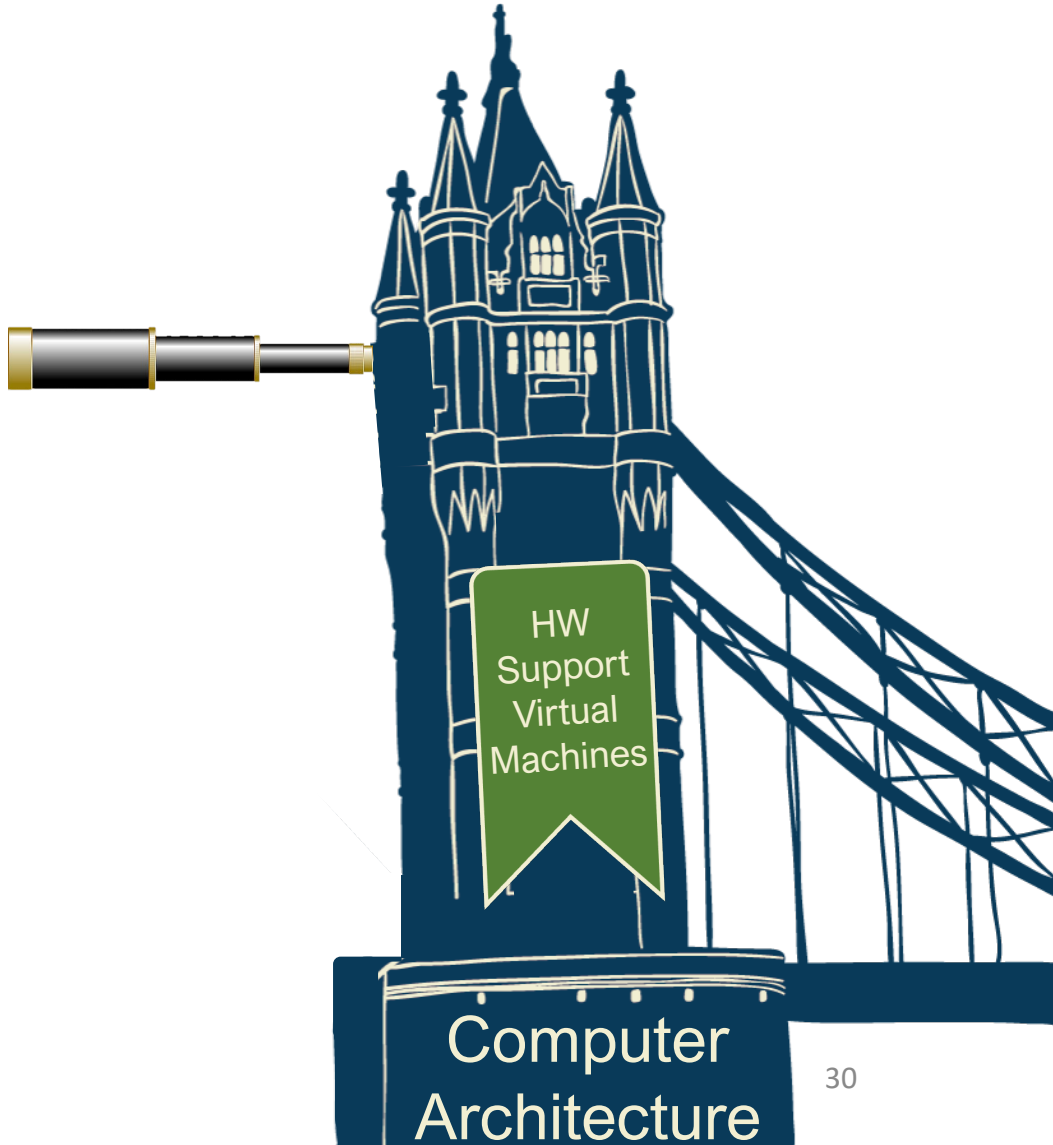

Reactive Research

Moore's Law

The opportunity

Rethink the Synergy between OS and HW

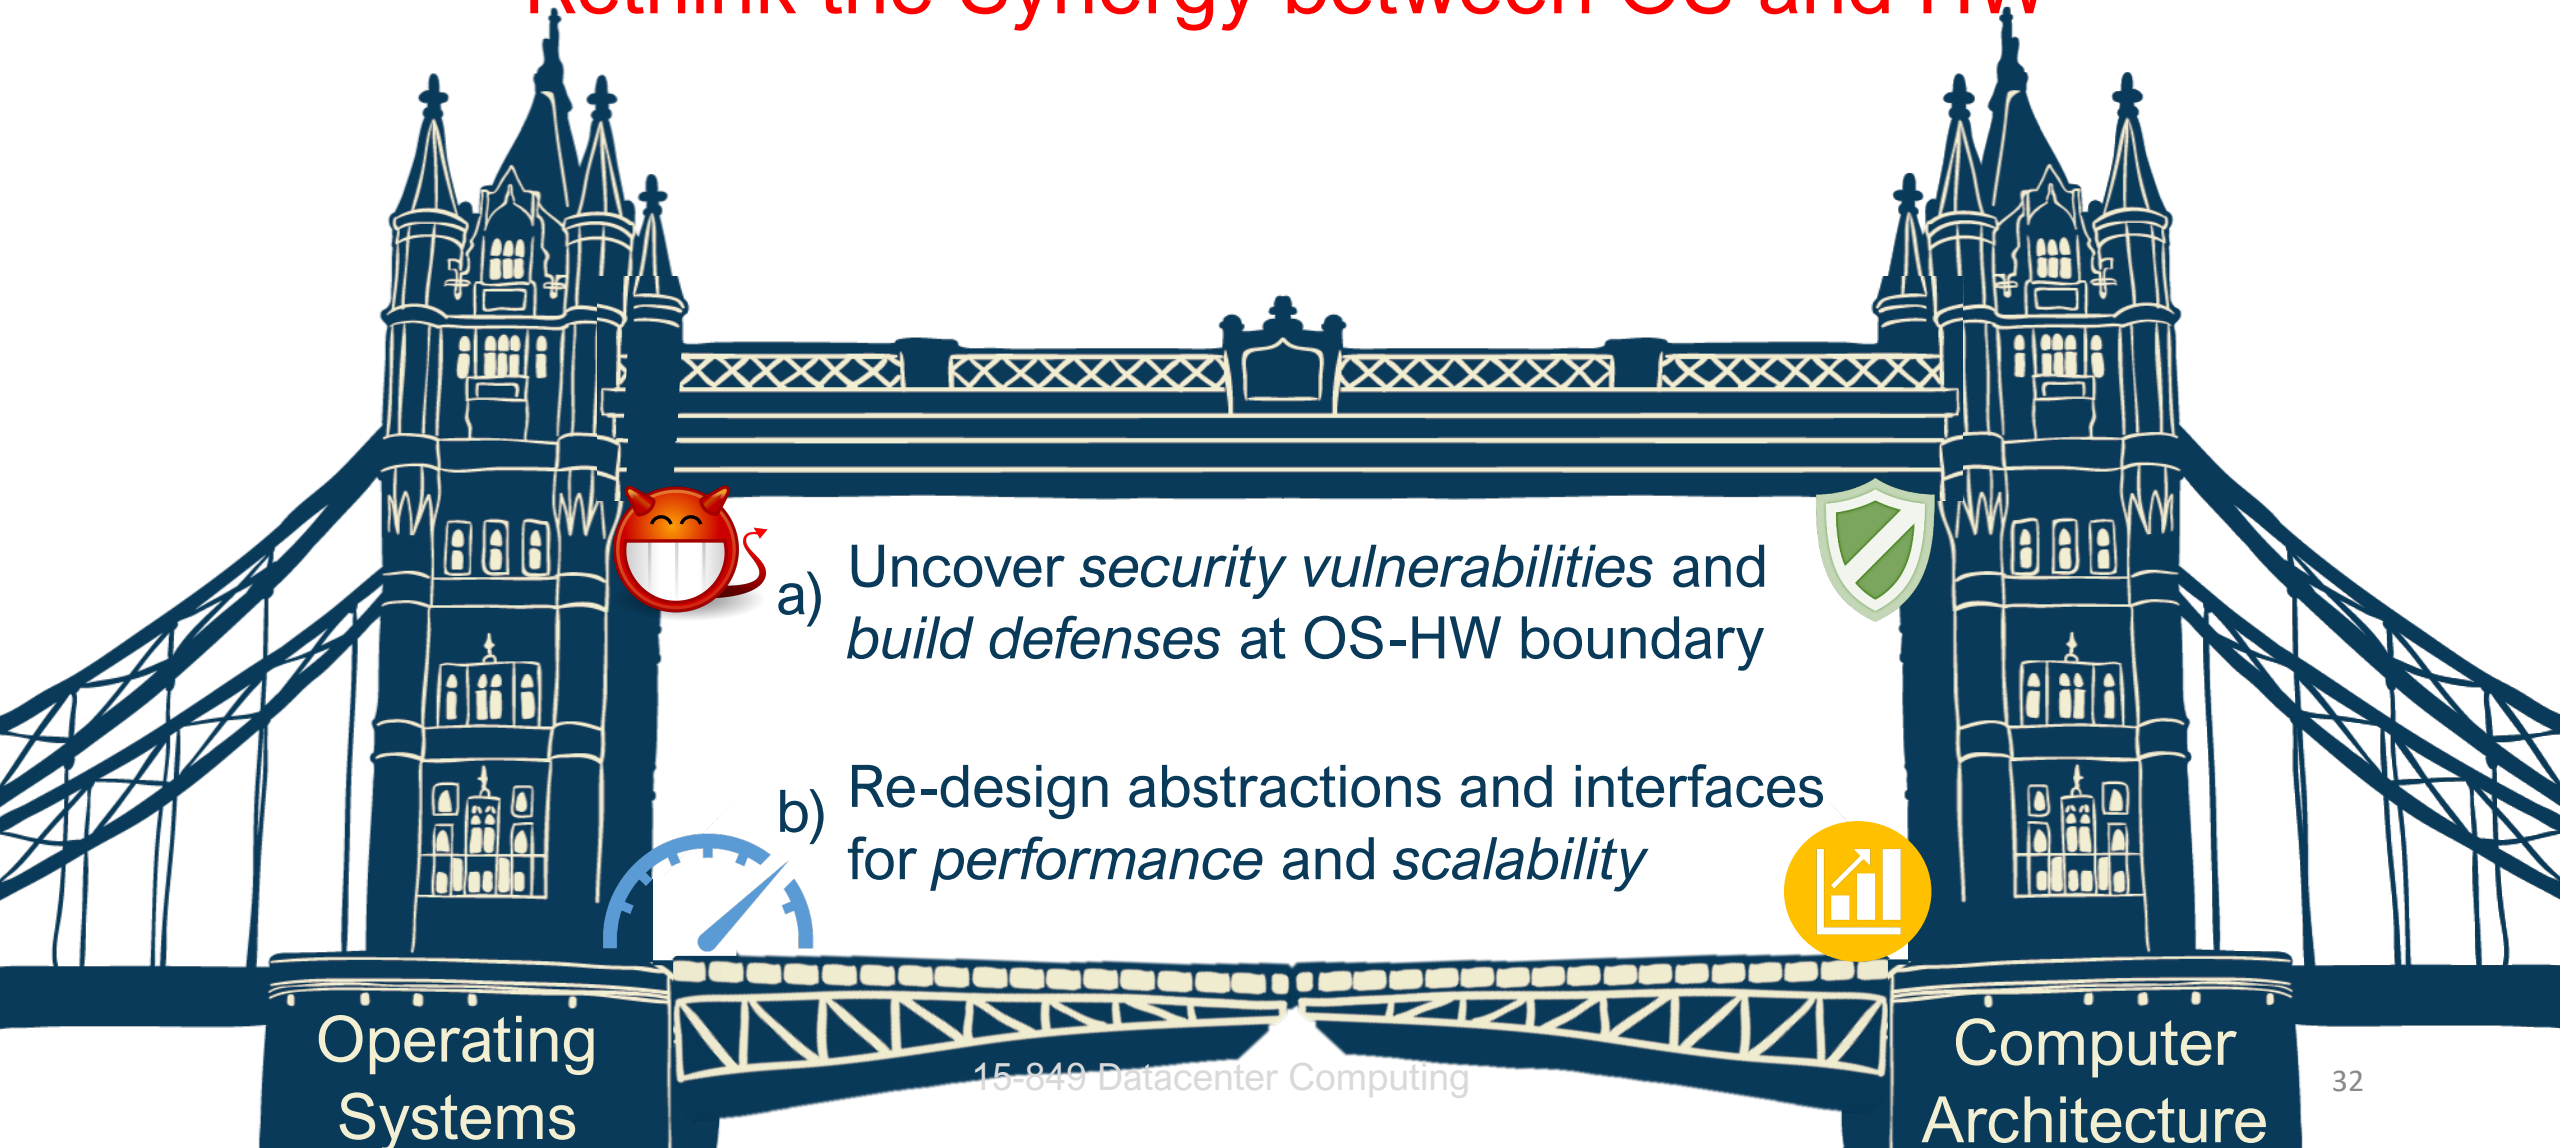

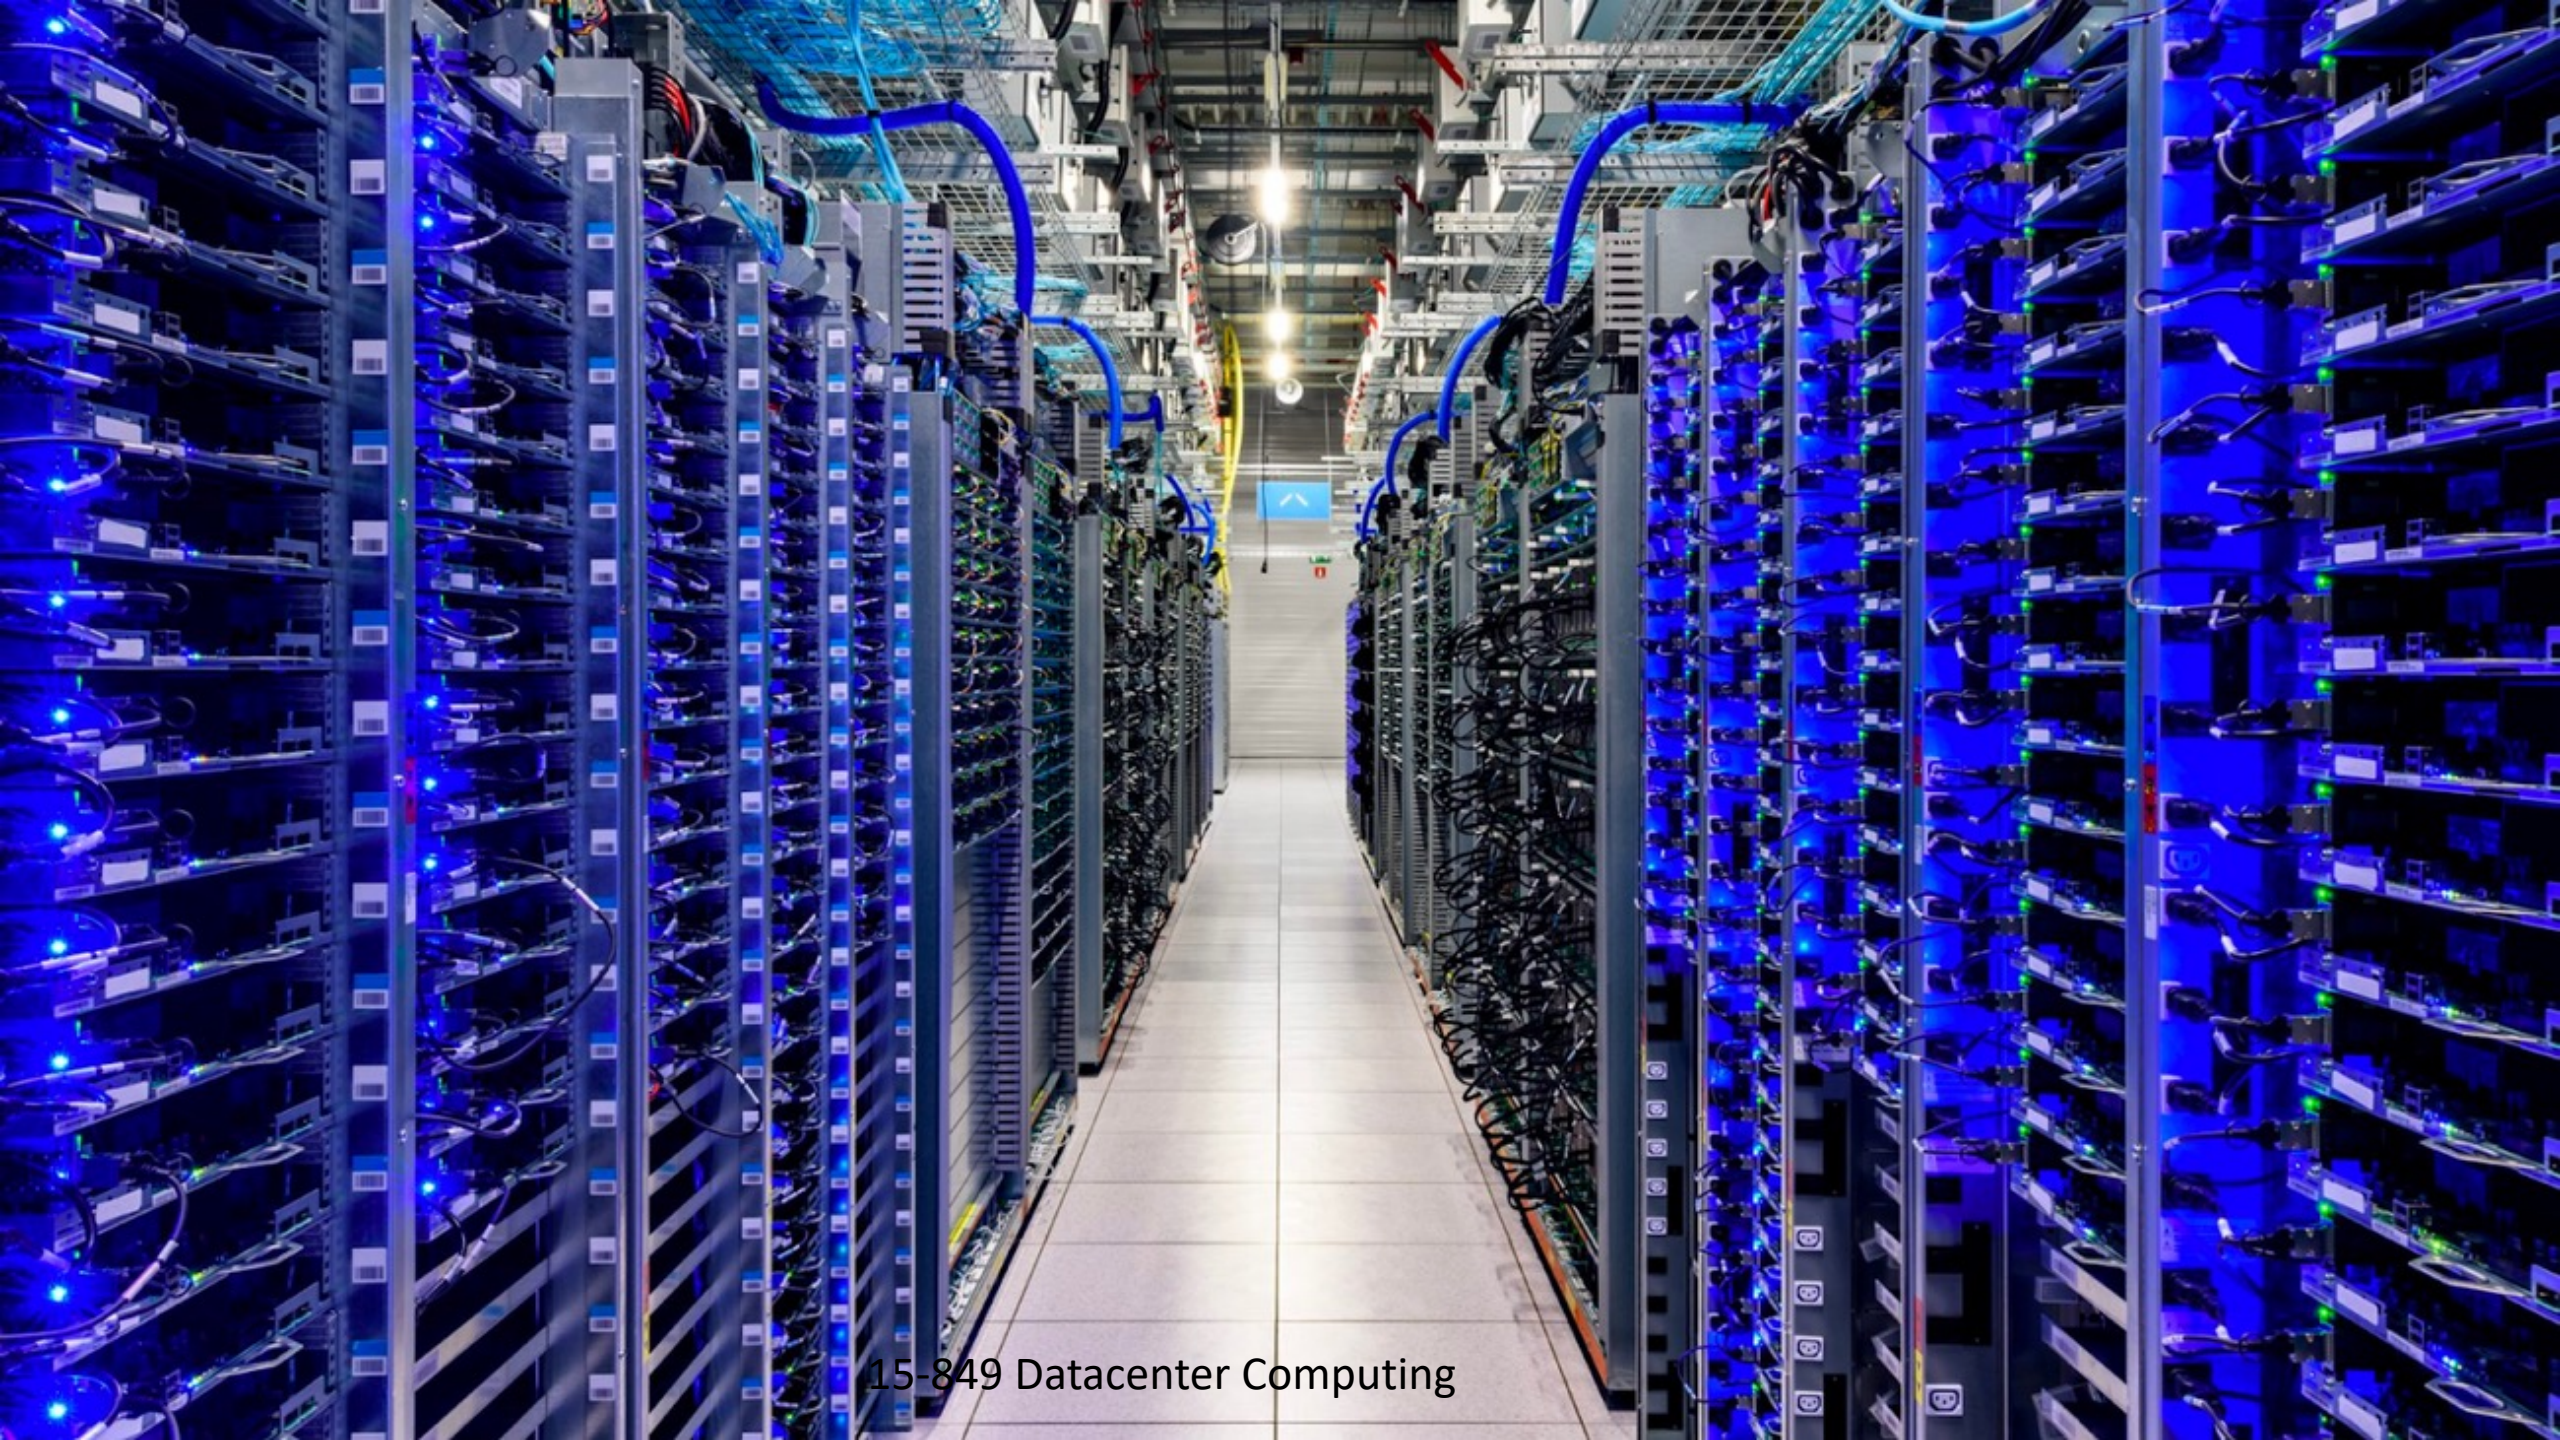

15-849 Datacenter Computing

Preview of Datacenter Topics

Side-channels in the Cloud
- Cores, Caches, Network

Transient Execution

Enclaves
- HW & Systems

Machine Learning Systems
- Industry talk

Memory Management
- Virtual memory
- Virtualization
- ML allocation

Accelerators
- TPUs and FPGAs

Warehouse-scale Analysis

Virtual Machines and Containers

Operating Systems & Low latency

Cluster Management
- Industry talk

Storage

Microservices and Serverless

Networking and Power Management

Why do Research Across HW and OS?

Software Features

Hardware Features

Reactive Features

"Please don't steal this totally wireless, magnetically mounted camera"

ars technica [1]

[1] https://arstechnica.com/gadgets/2021/08/new-google-nest-cams-can-record-video-without-a-monthly-subscription/

Why Hardware Security?

End goal:
- Shared hardware
- Resource virtualization
- Secure Datacenter Computing!

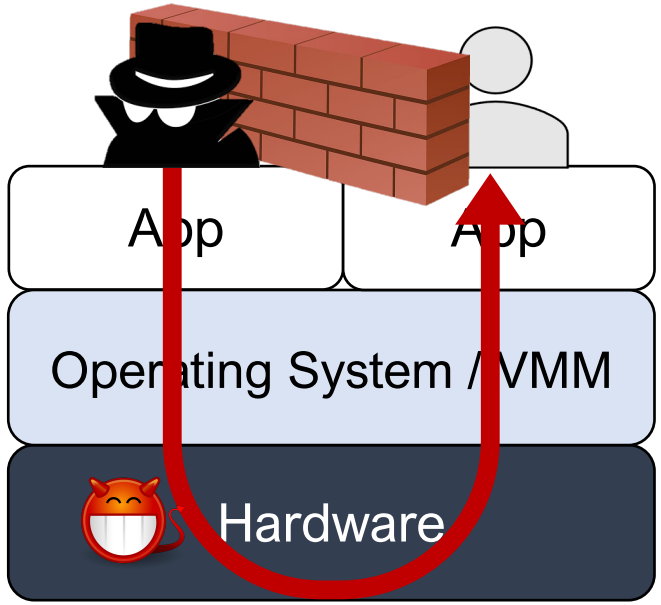

TCB: Trusted Computing Base
- SW + HW

Shrinkage in SW :
- OS
- Hypervisor

HW remained almost the same

Advanced over time:
1980s → IBM mainframes (temp sensors for secure boot)
2000s → ARM TrustZone, Intel TPM
2010s → Intel SGX, AMD SEV-SNP

The Era of Side-Channels

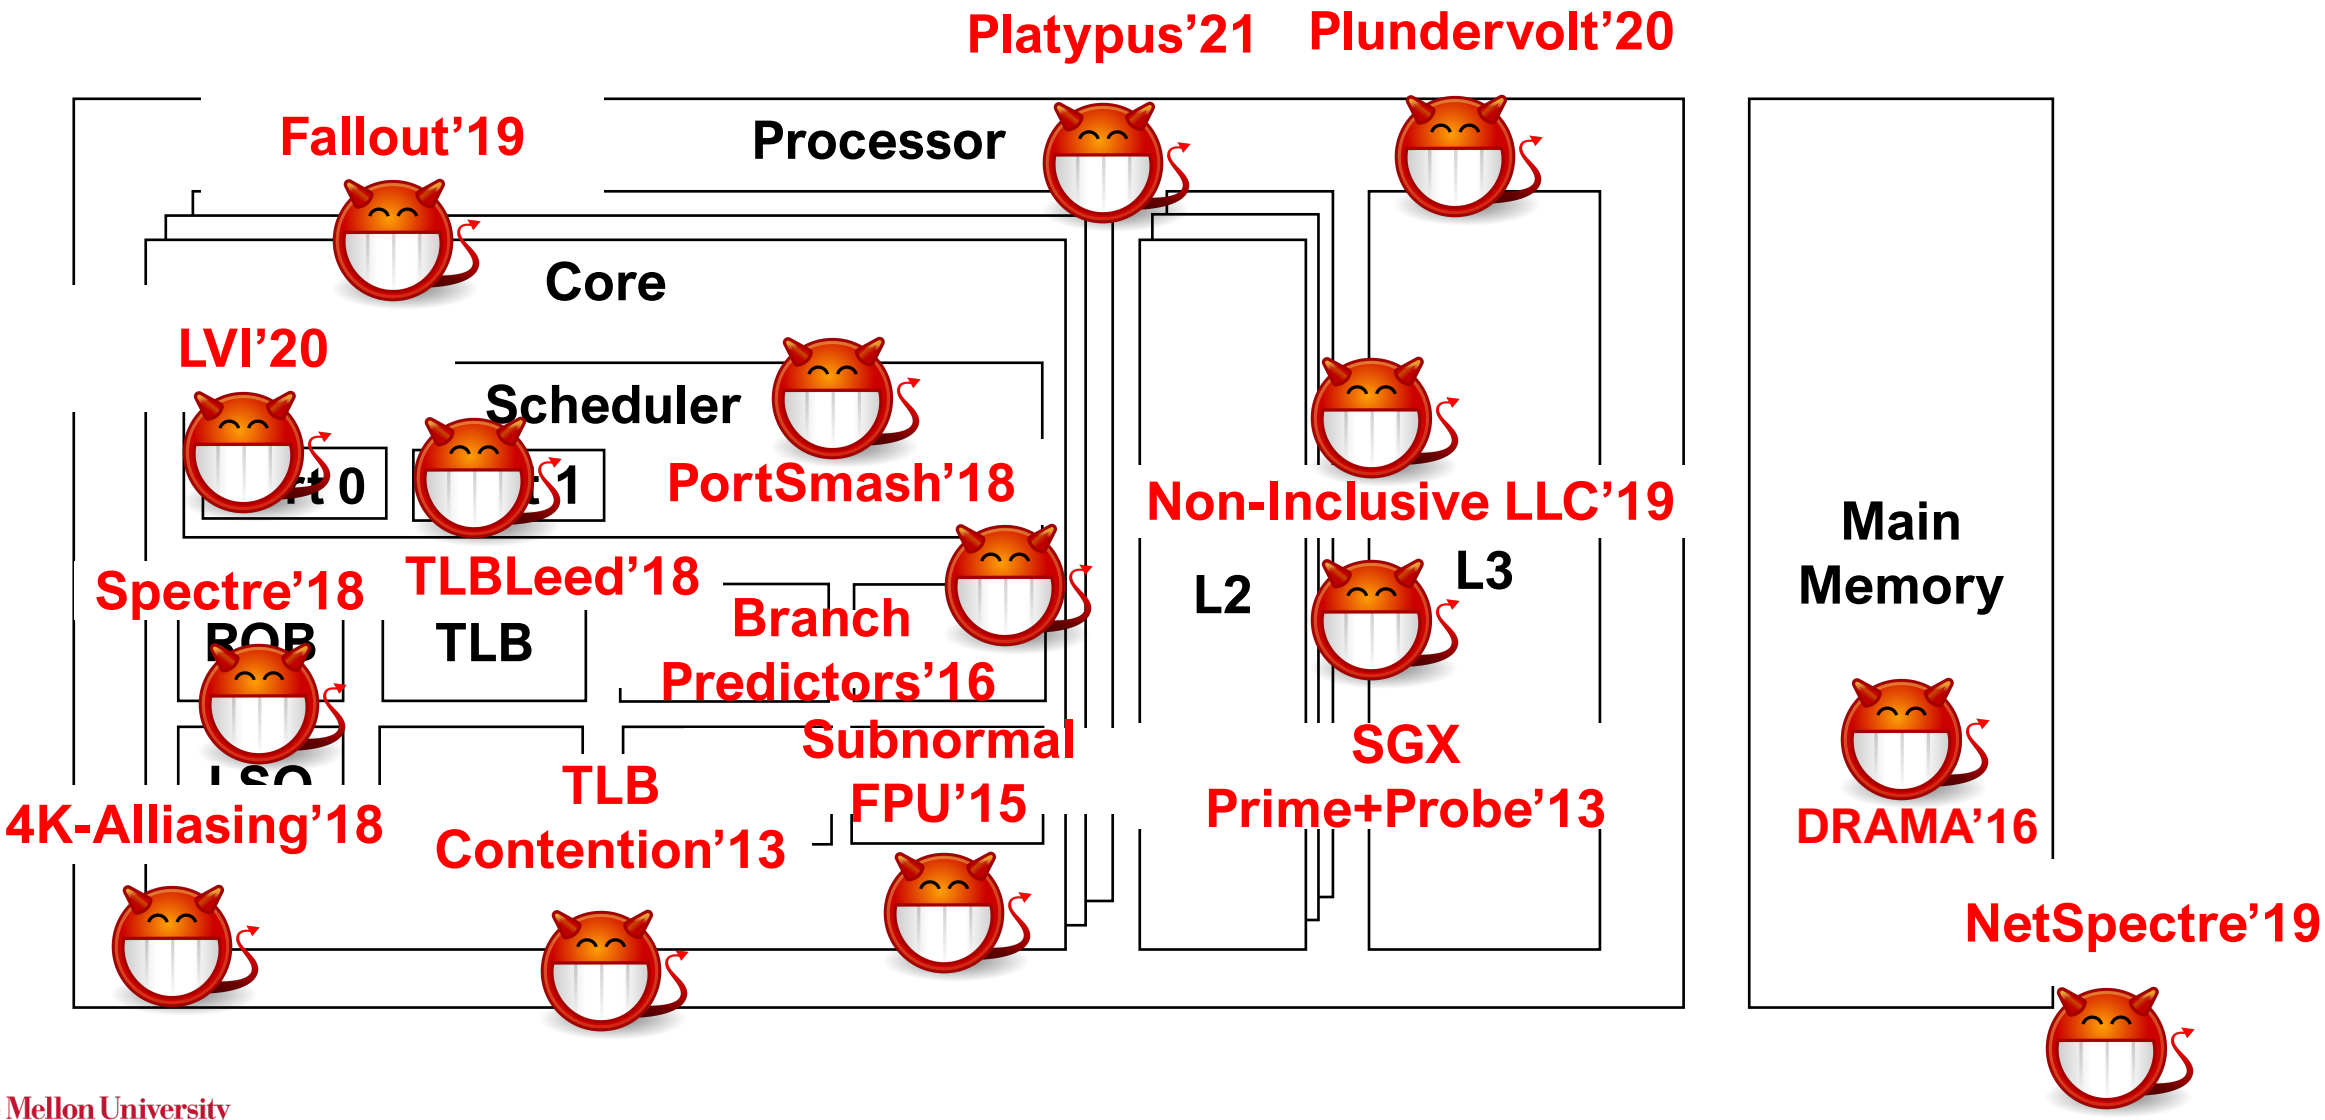

Microarchitectural Side-channels

Secret-dependent Execution { Array[secret] }

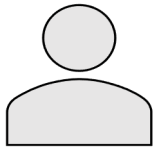

The Channel
Any micro-architectural Structure

{ Transient/Non-transient } X { Core, TLB, Caches, DRAM }

Microarchitectural Side-channels

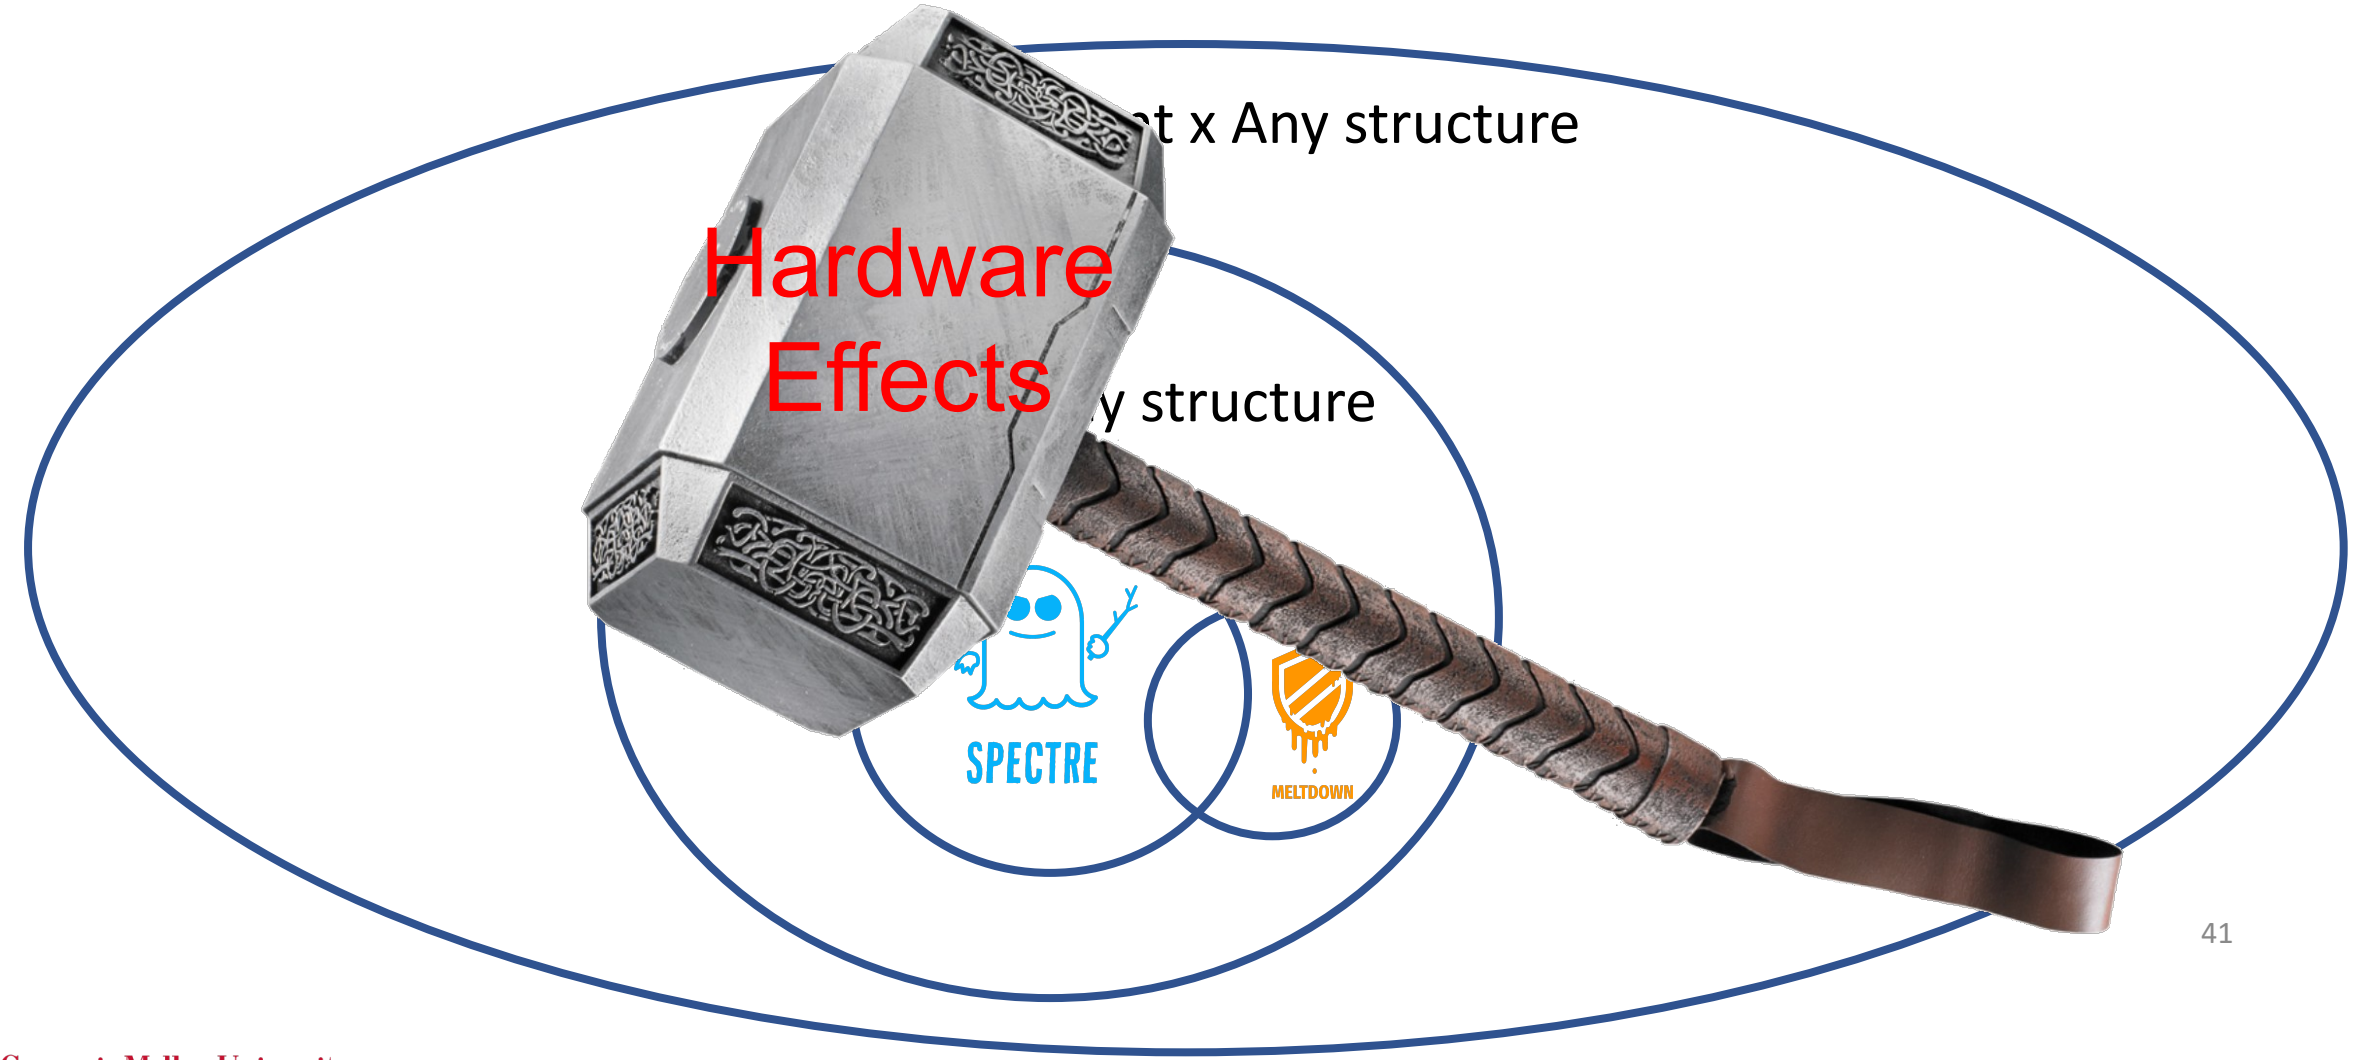

Enclaves

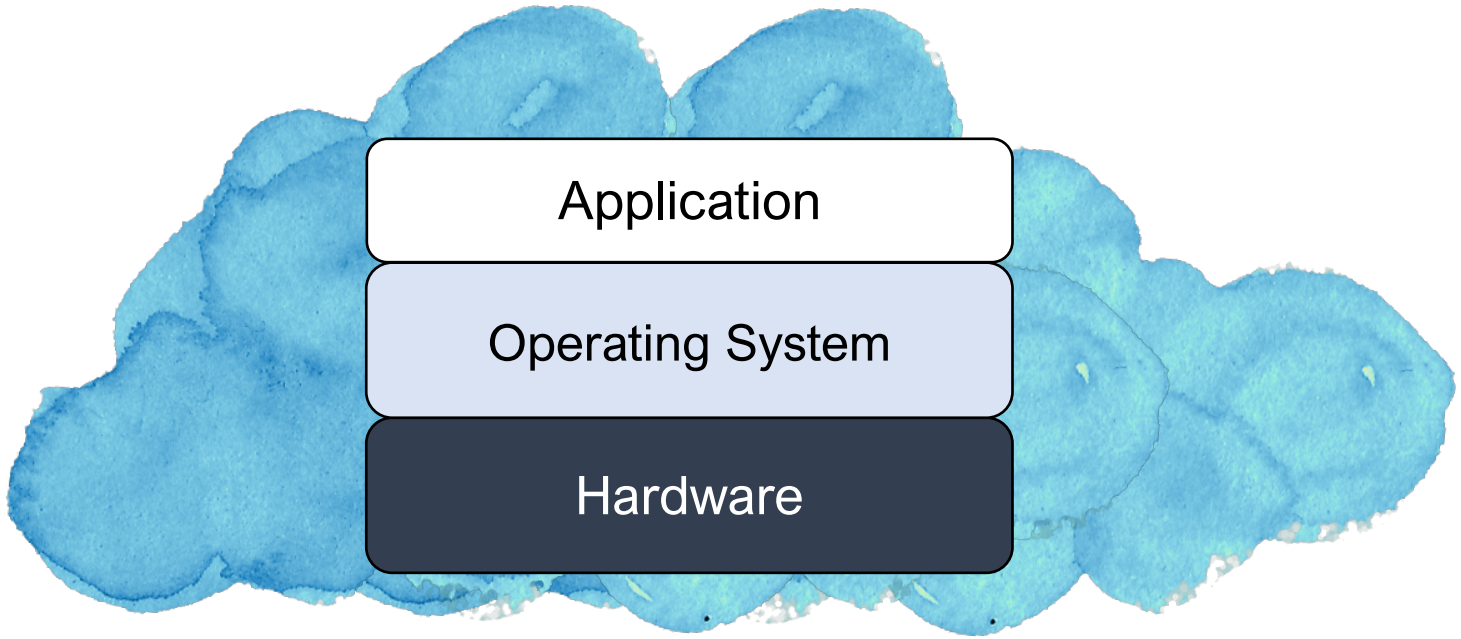

Enclaves

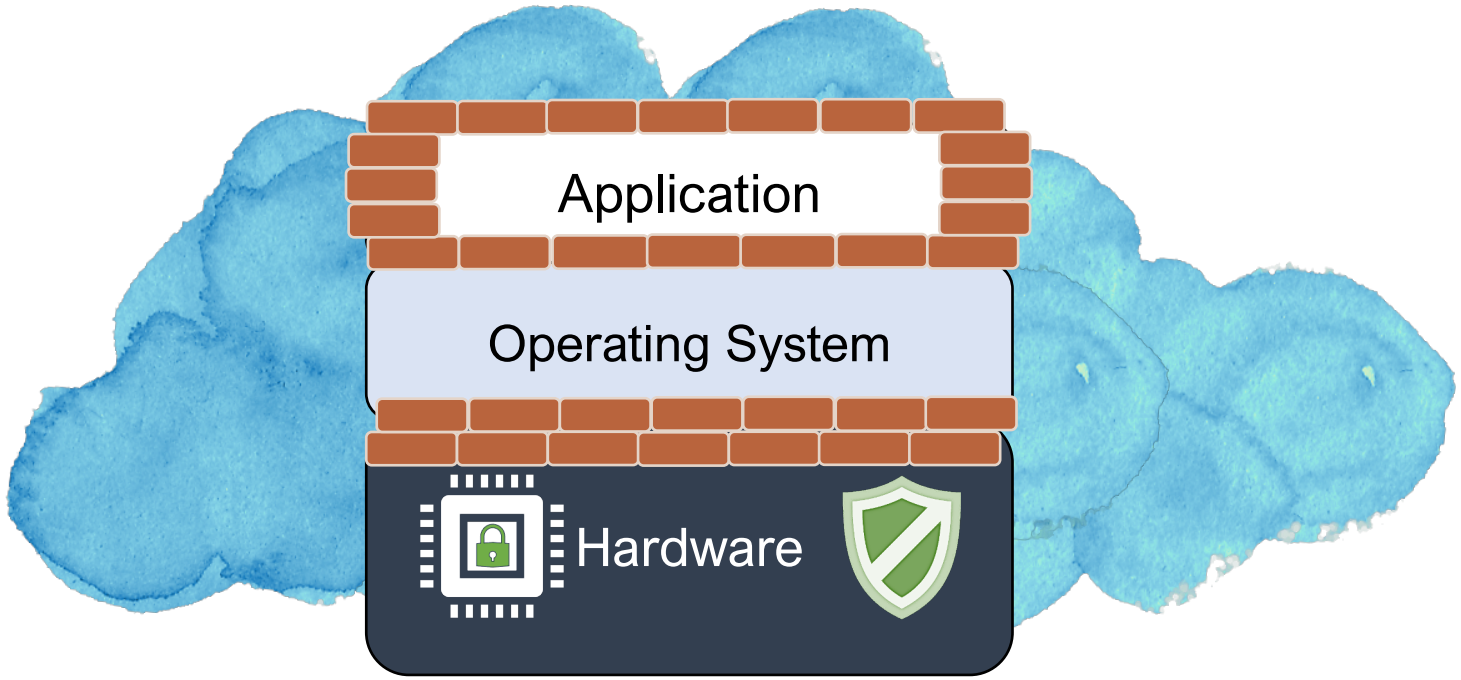

Enclaves

Do not trust OS/Hypervisor

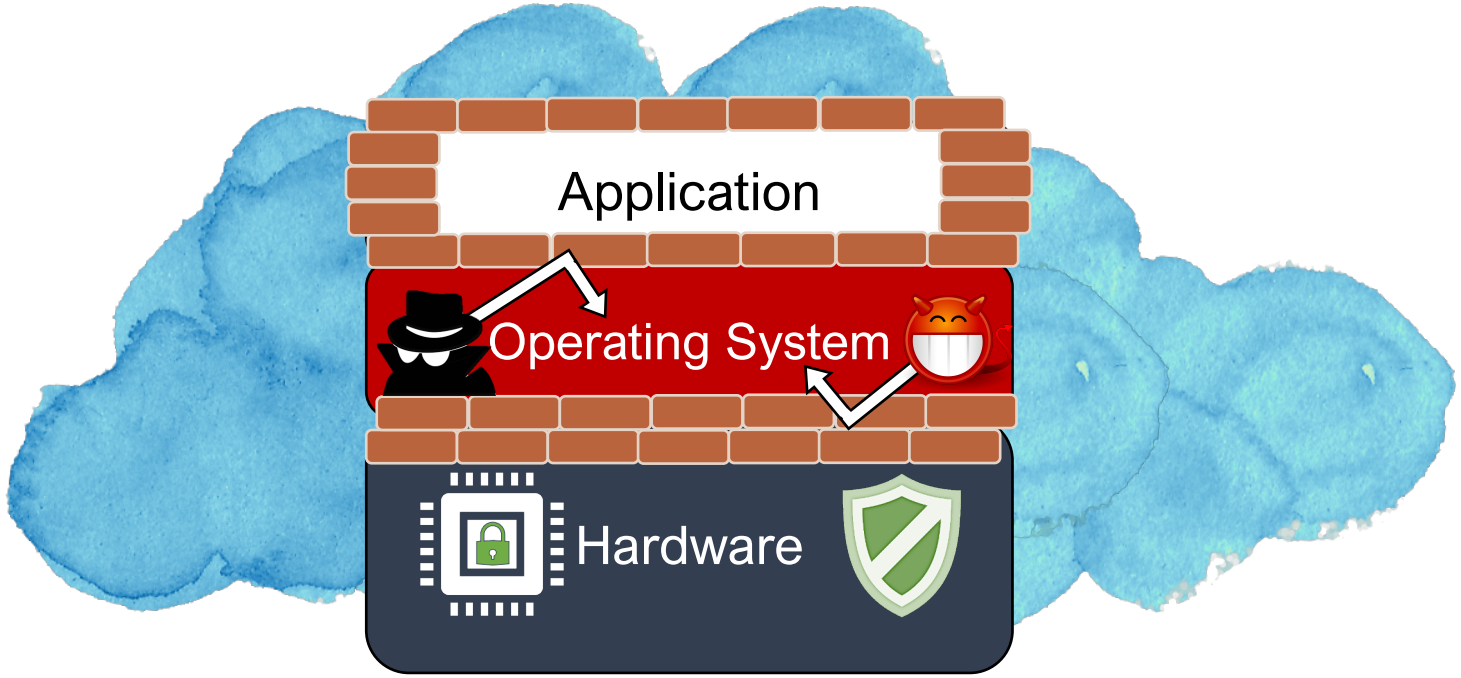

OS/Hypervisor still performs some management operations!

Enclave Design

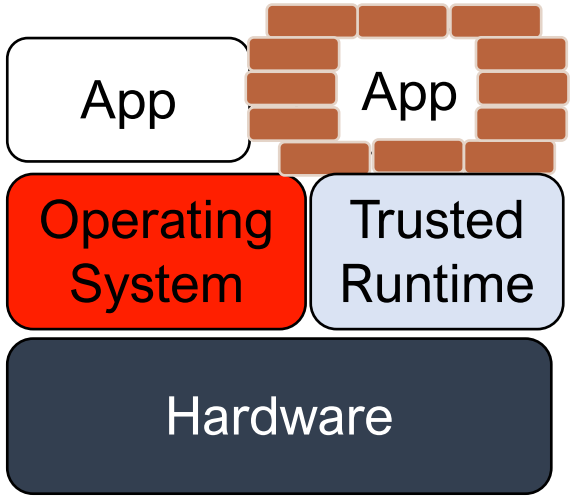

- How do we design enclaves?
- How are they integrated in the System?
- What about side-channels?
- Can we really program with enclaves?

The Memory Problems!

The Memory Problems!

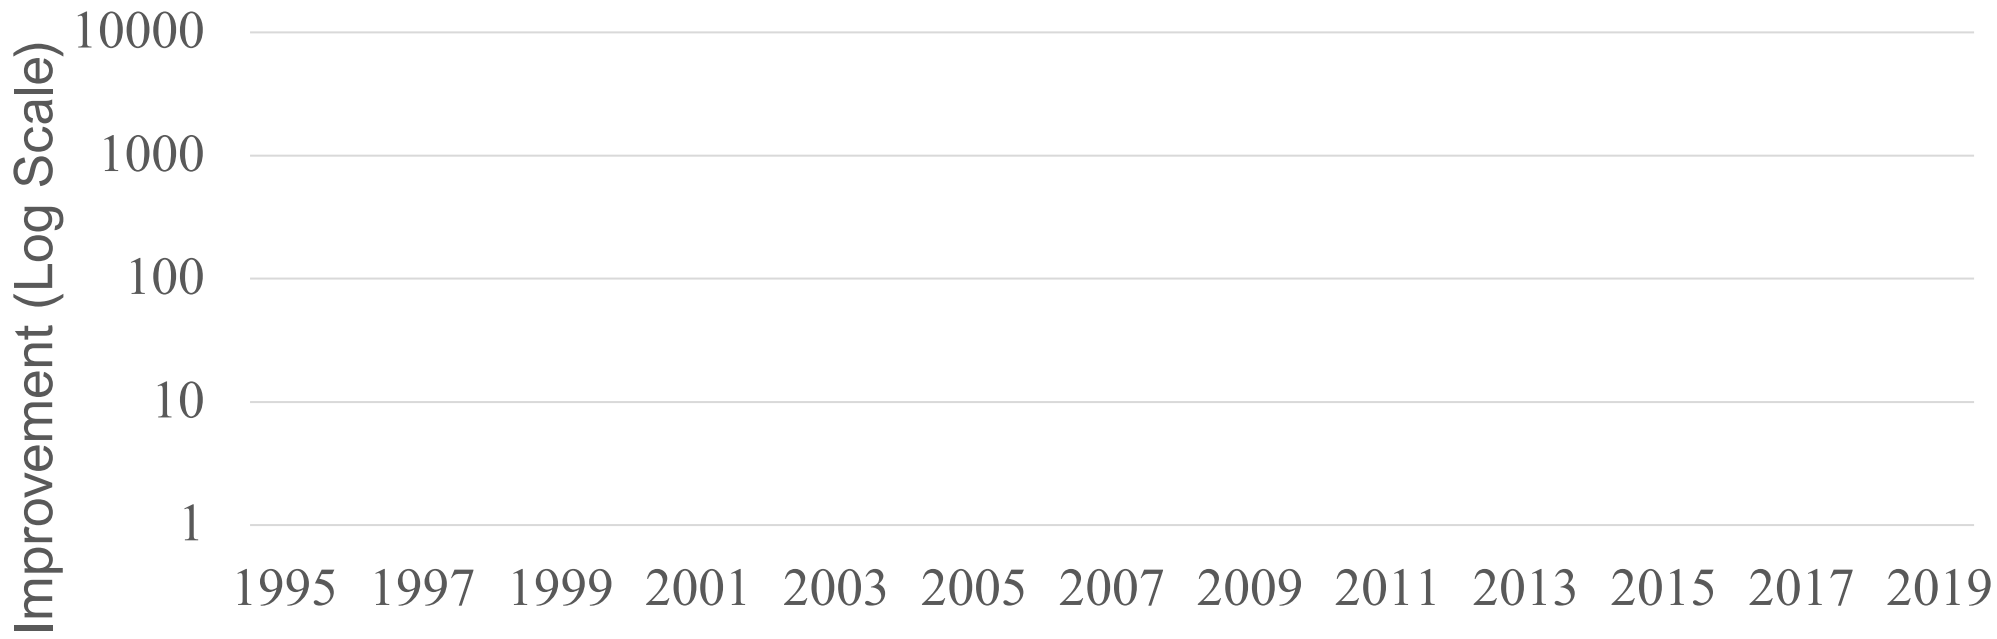

The Memory Problems!

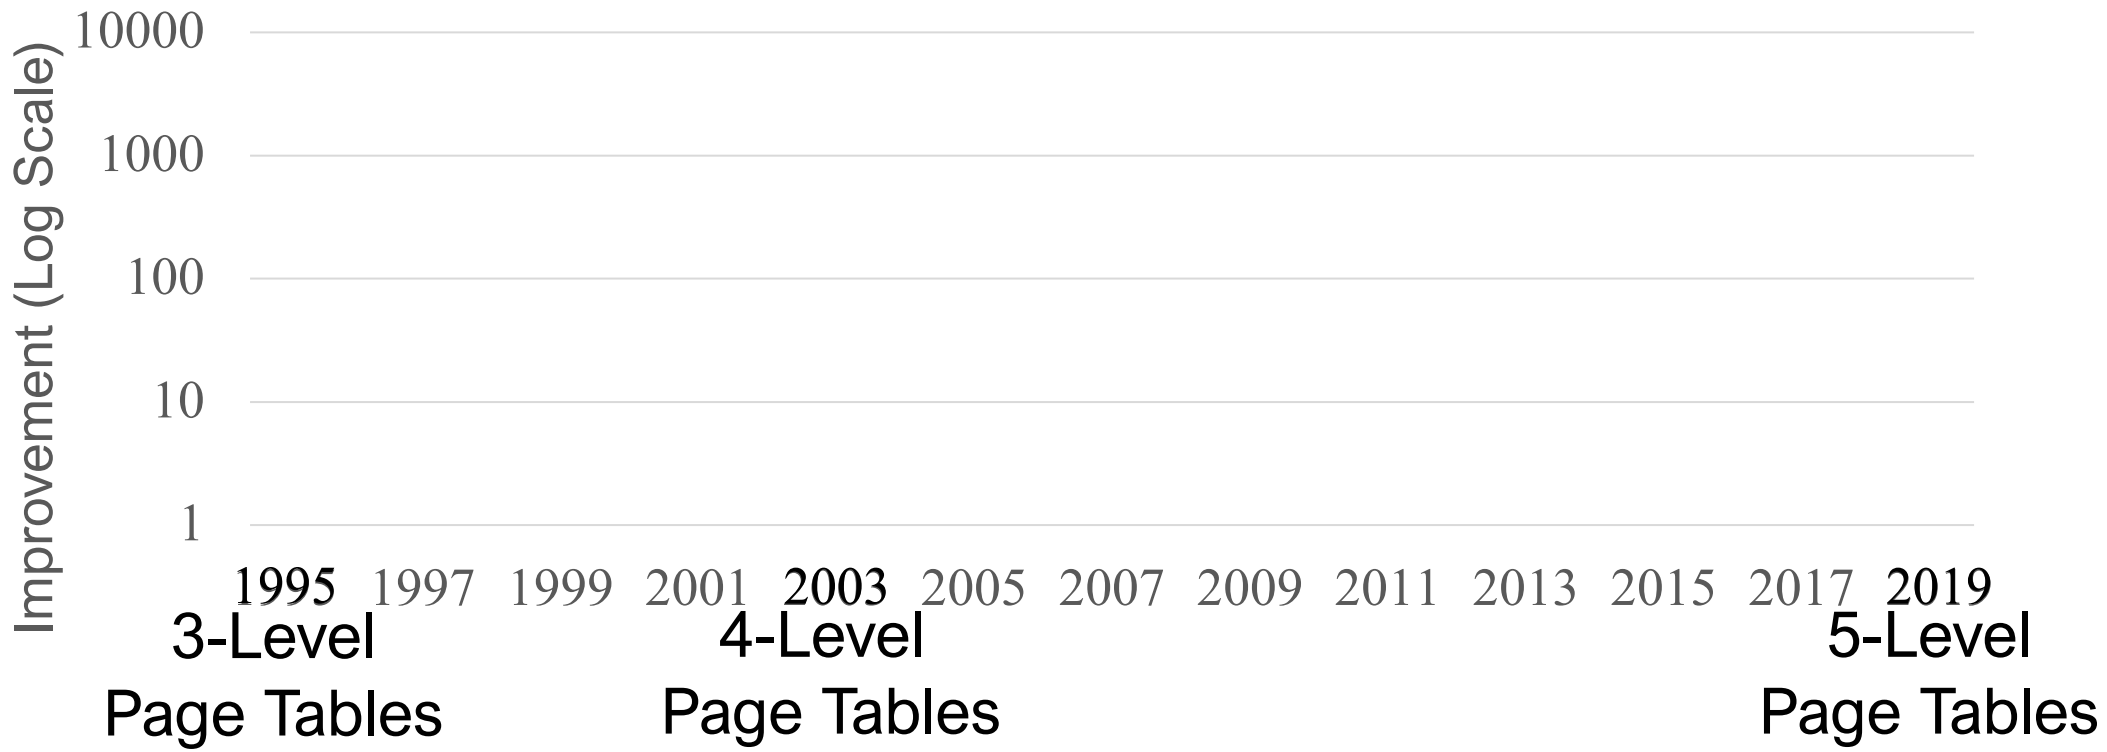

Time to Rethink Memory Management

Current Memory Management is not Scalable!

Virtual Memory

ML-guided Memory Allocation

Memory Virtualization

Heterogenous Datacenters

The opportunity of accelerators

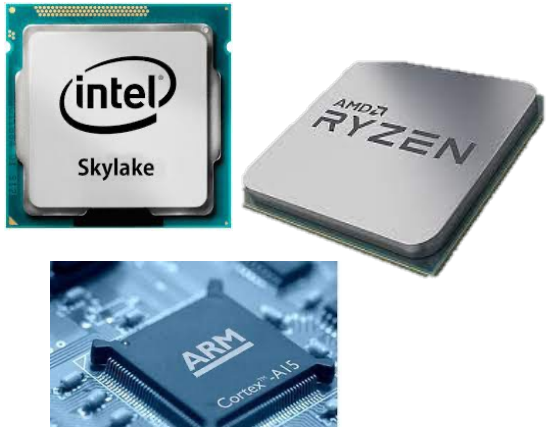

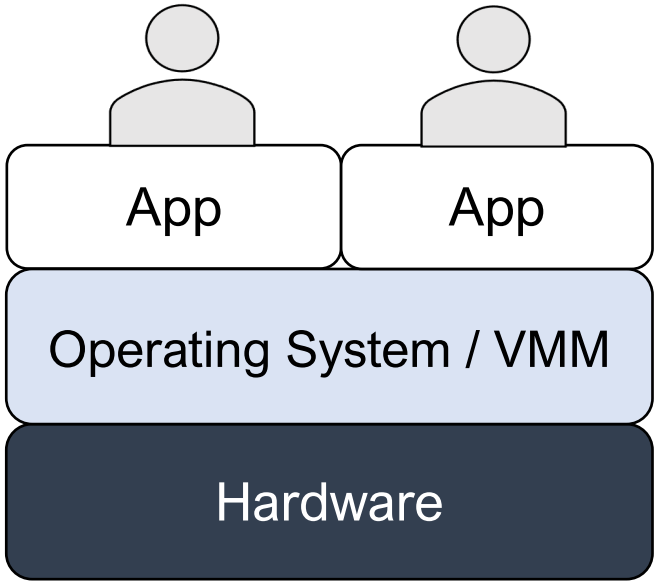

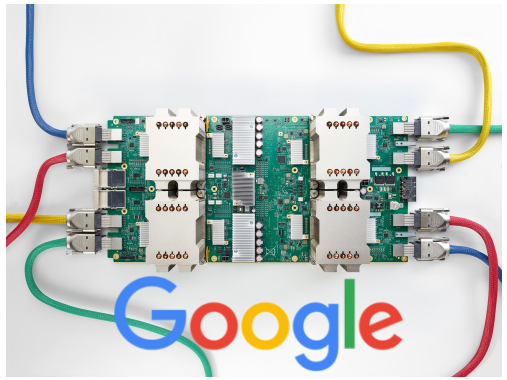

Lessons learned from
5 years of accelerator research

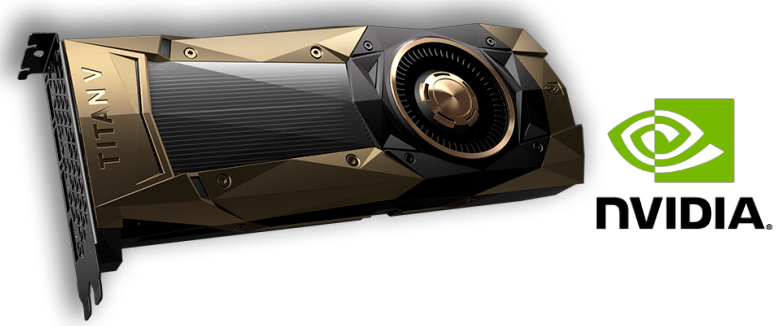

Containers and Lightweight Virtualization

Conventional Datacenter Computing

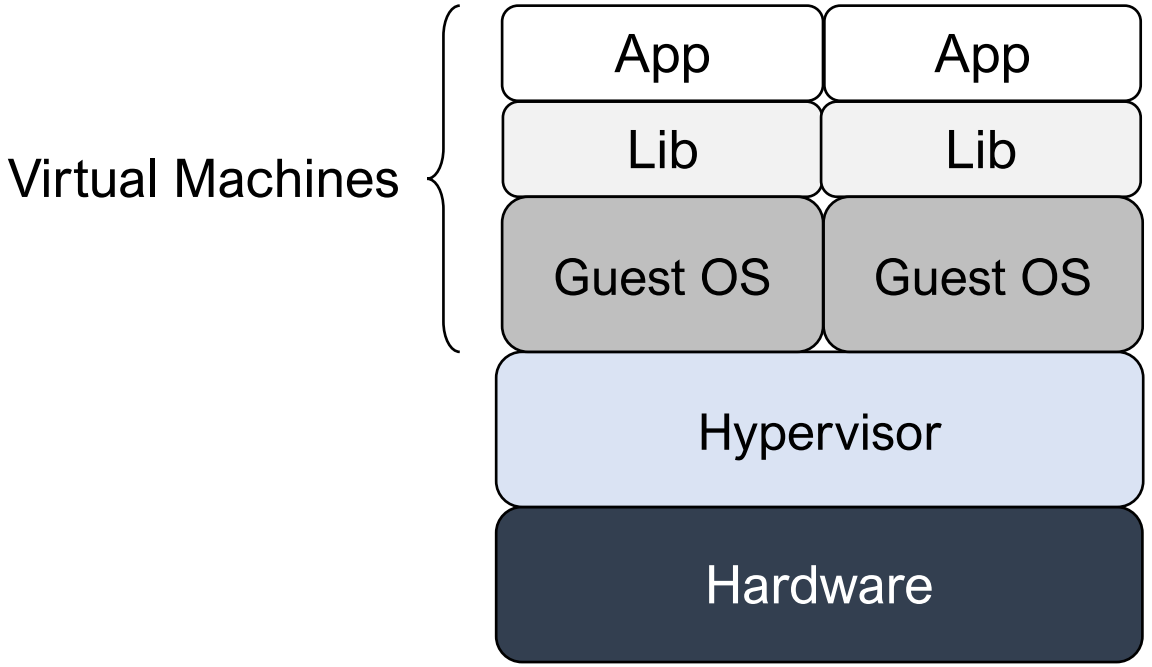

New Era in Datacenter Computing

1. Lightweight
2. Faster bringup
3. Higher consolidation

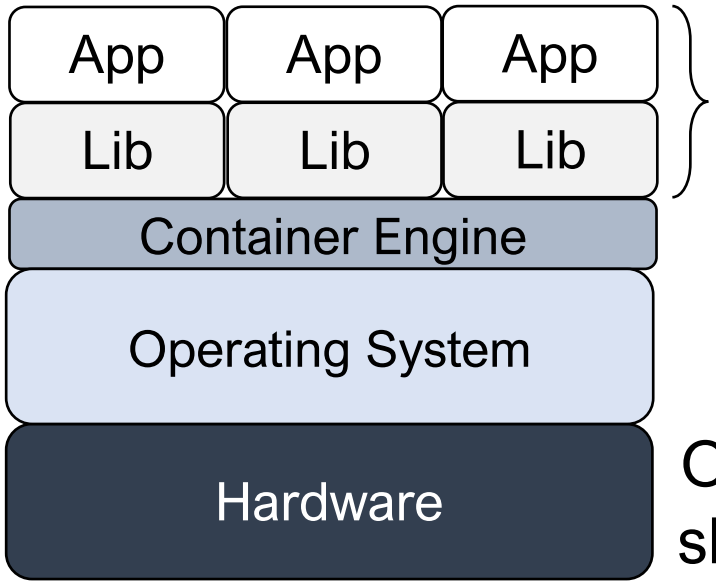

Containers

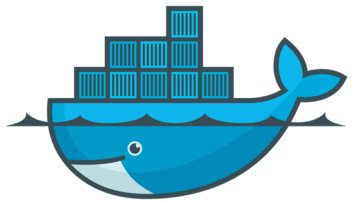

docker

OS automatically shares resources!

Containers in the Datacenter

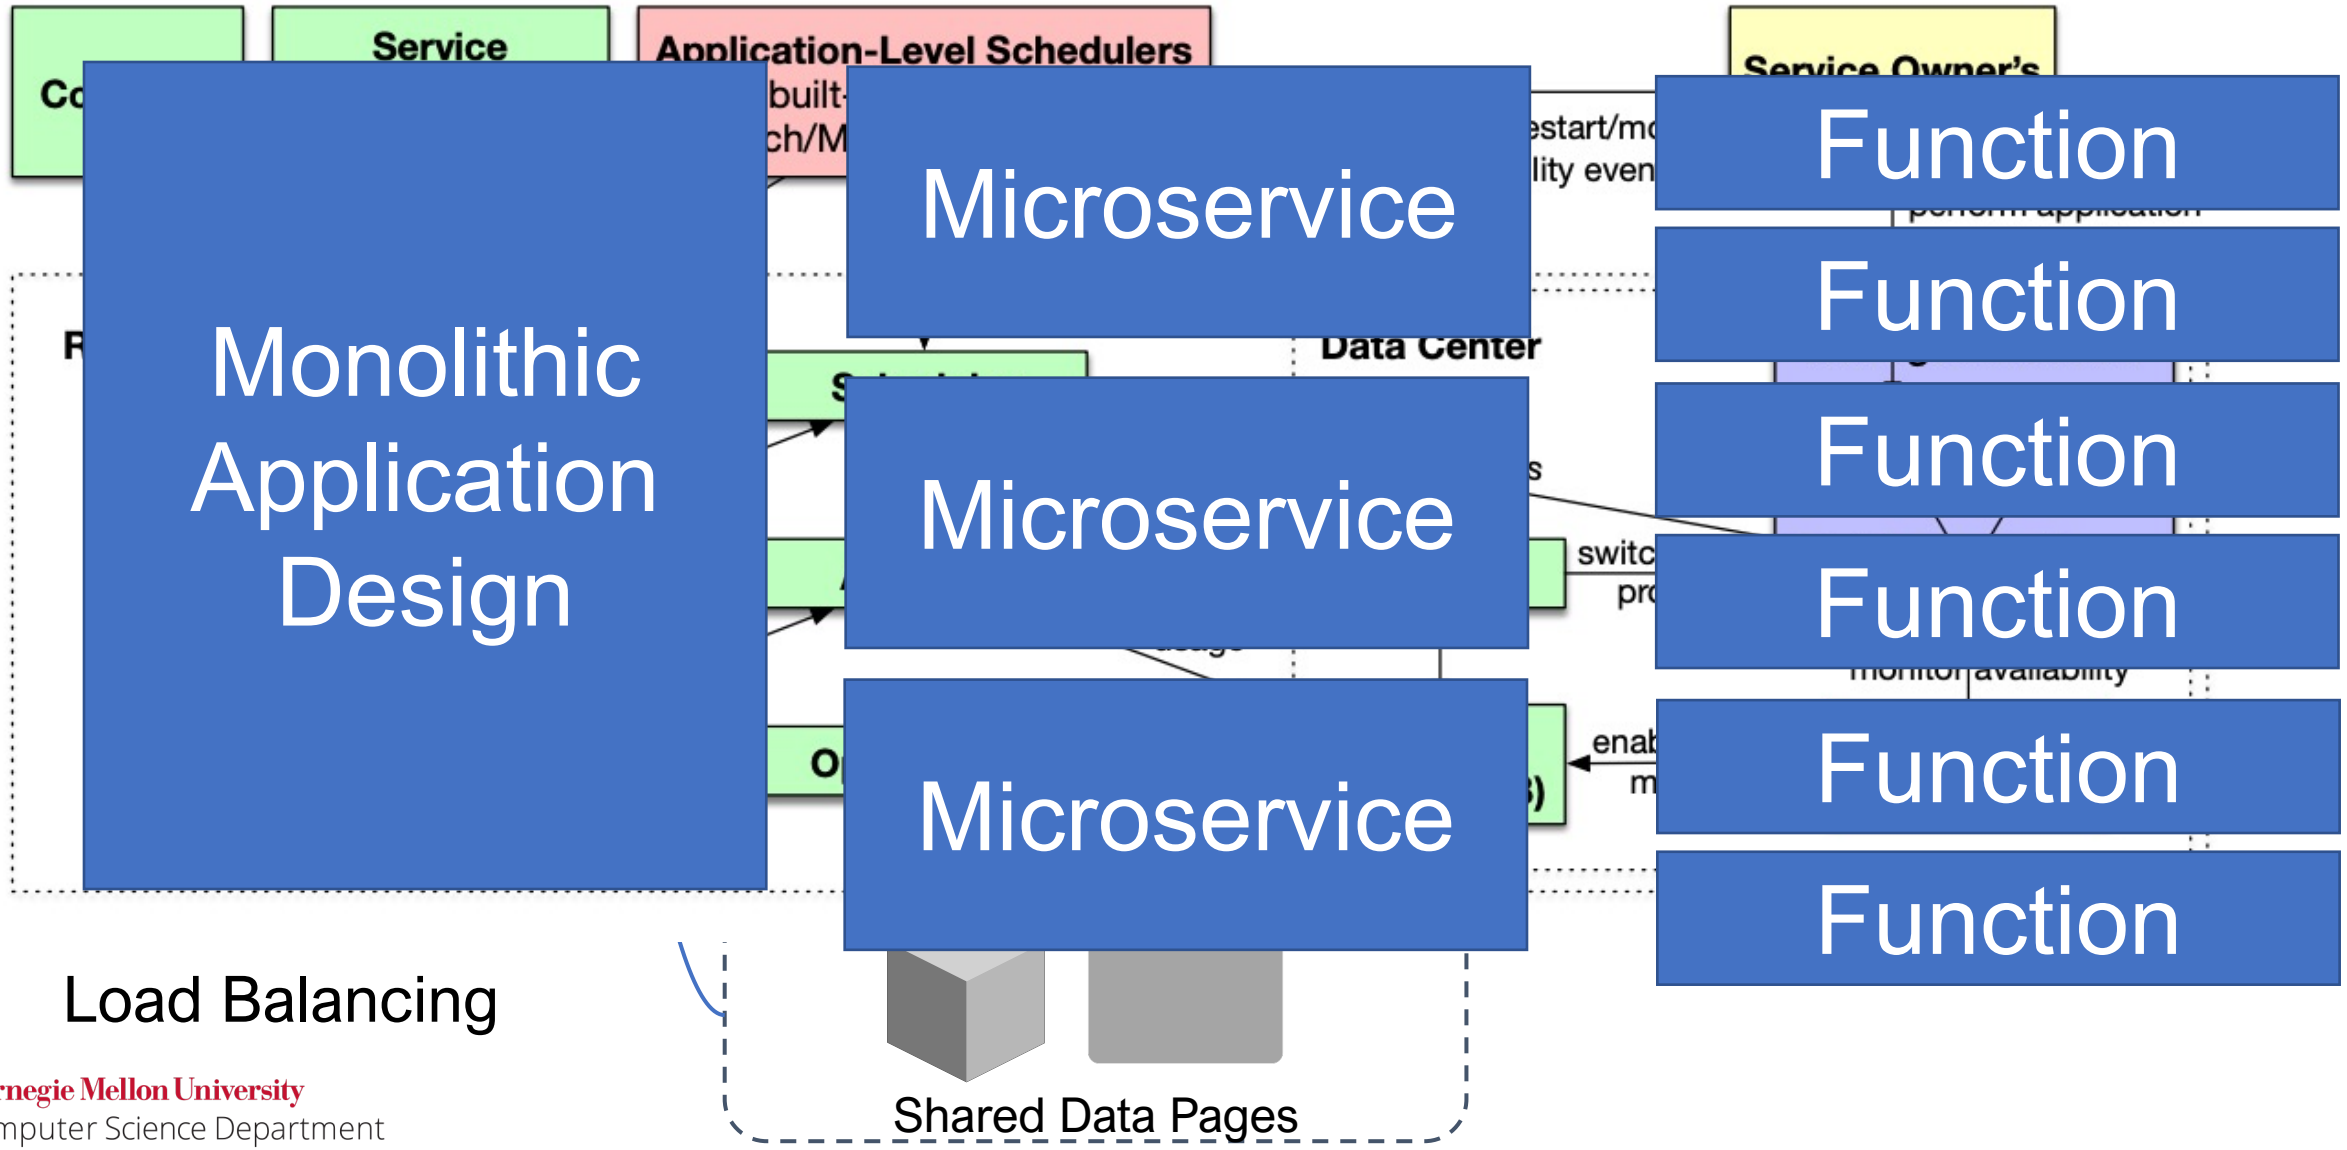

Security Limitations of Containers

1. Shared OS
2. Attack surfac

ers

CVE-2018-18281 | CVE-2017-5123 | CVE-2017-18344 | CVE-2016-5195

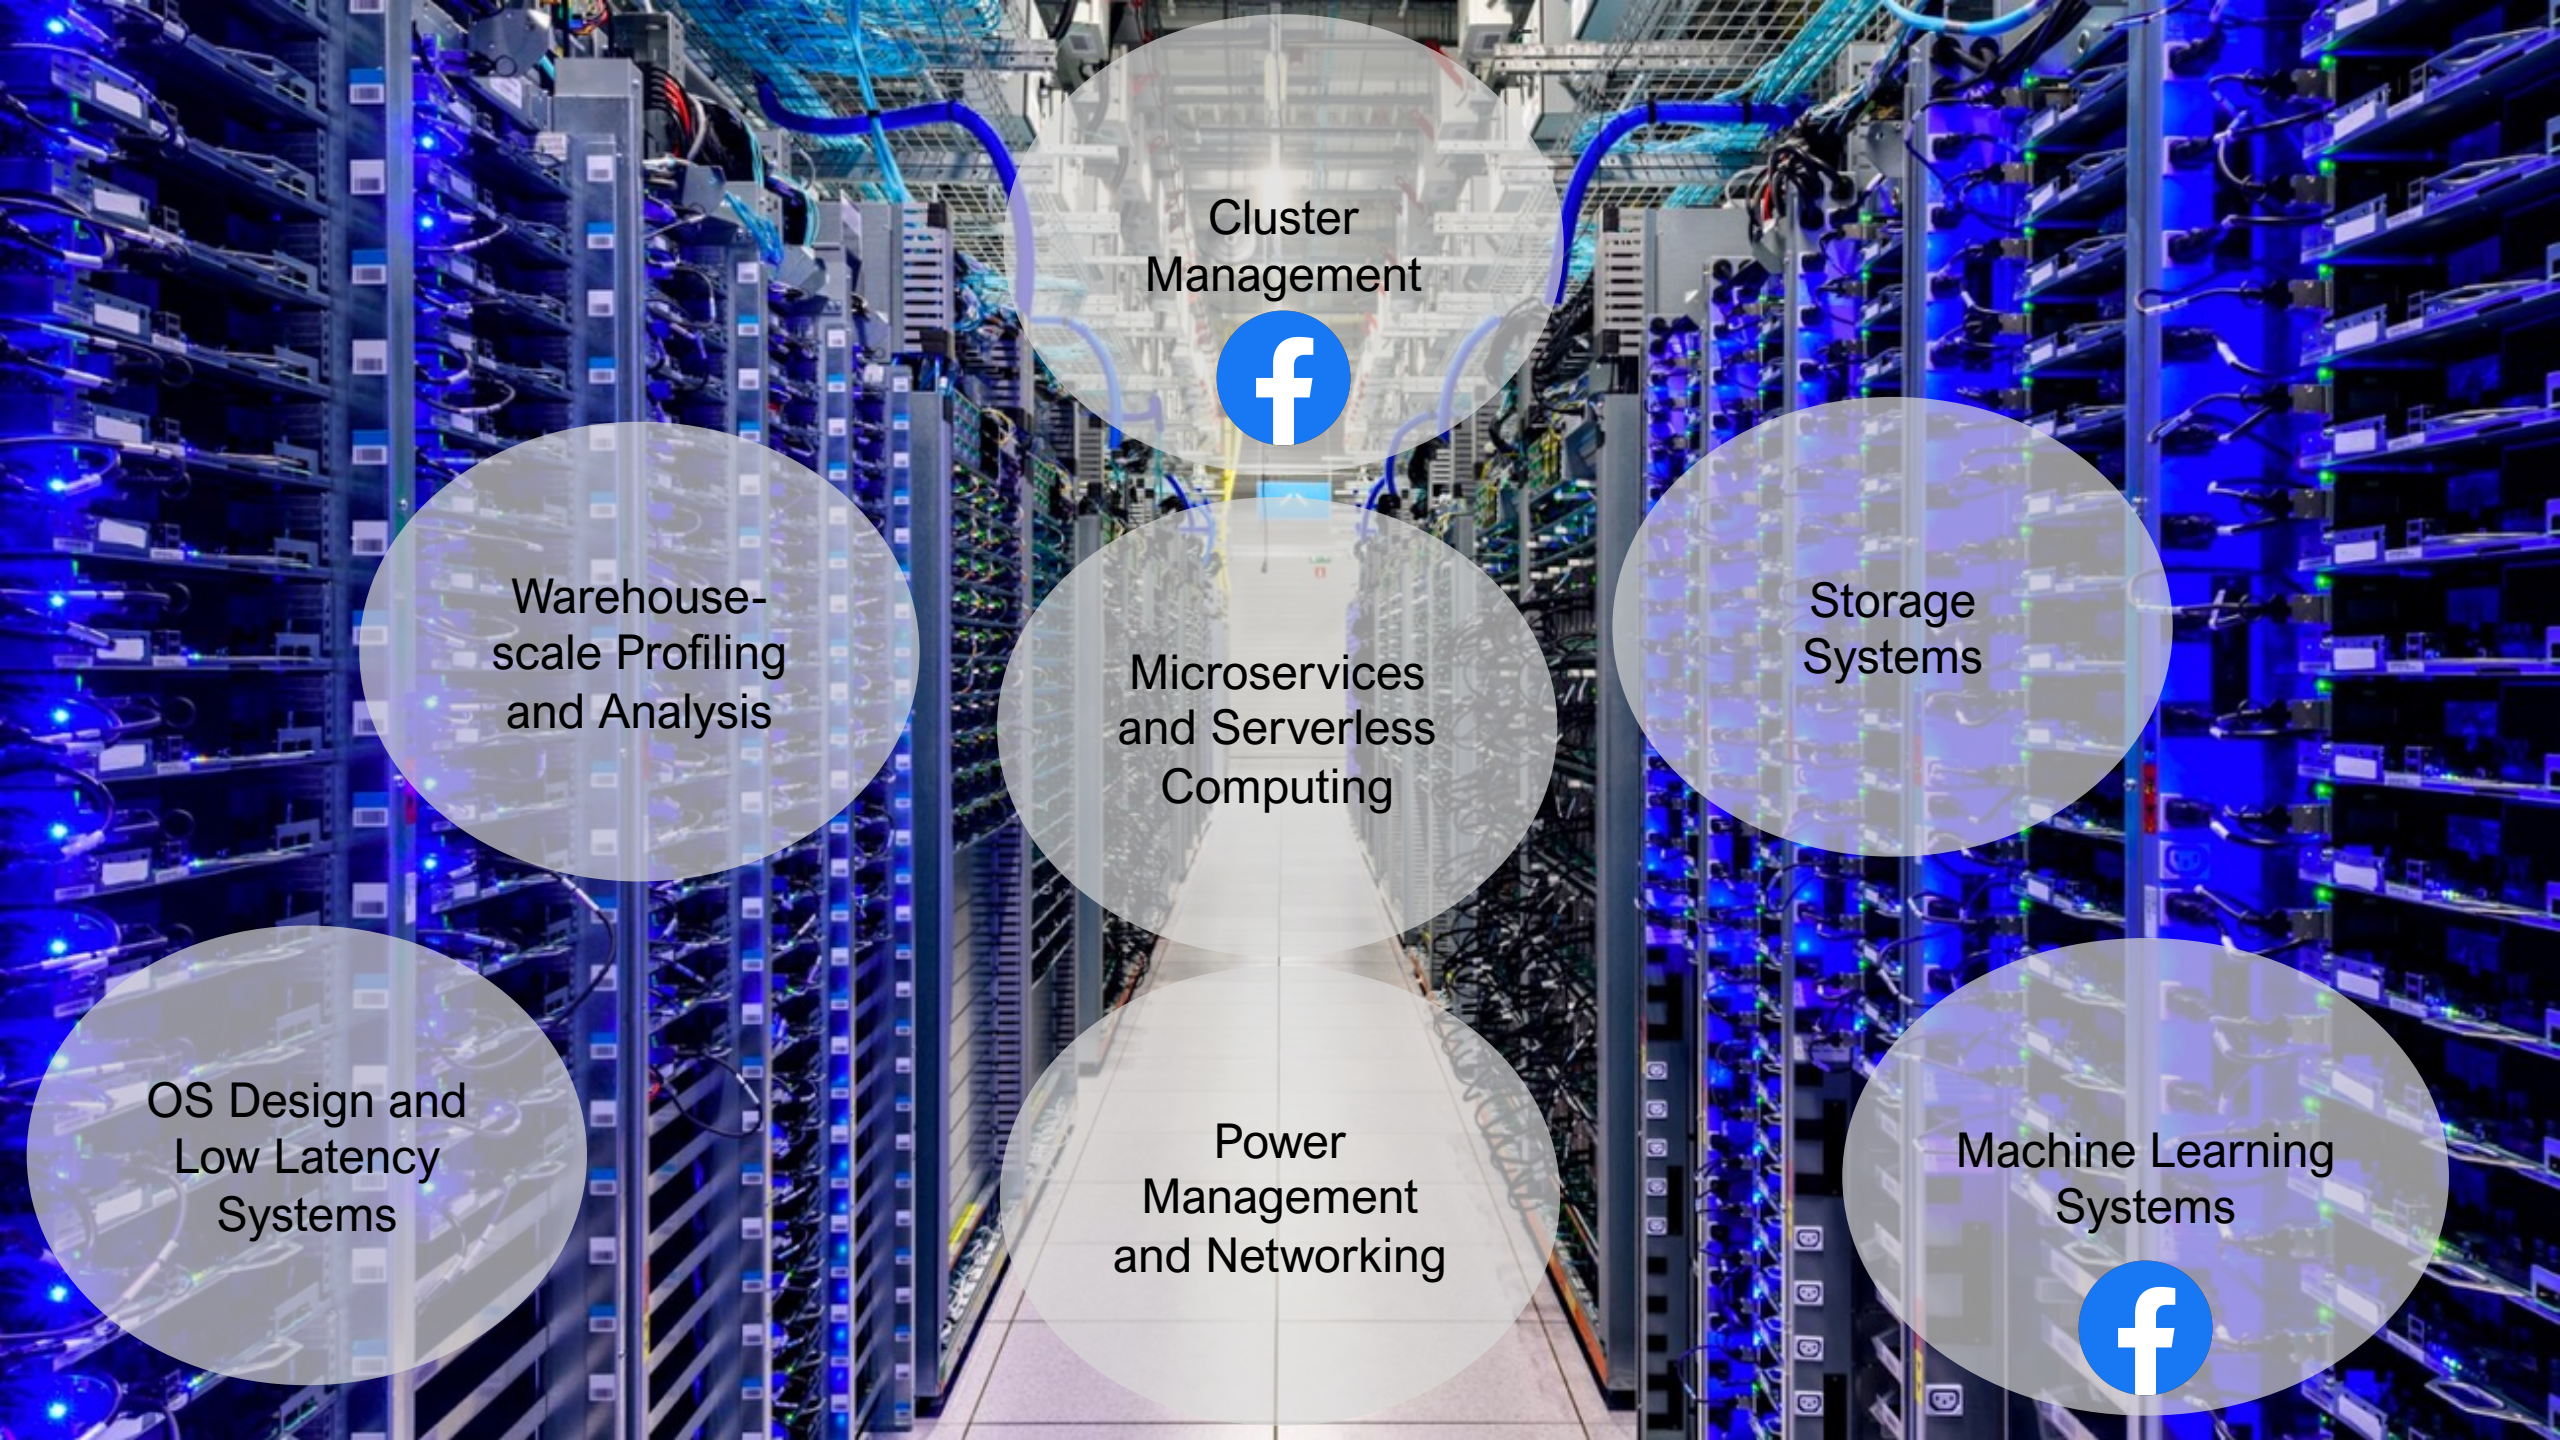
Cluster Management
Warehouse-scale Profiling and Analysis
Microservices and Serverless Computing
Storage Systems
OS Design and Low Latency Systems
Power Management and Networking
Machine Learning Systems

Next Up → Background in Architecture!

Check paper schedule

- https://www.cs.cmu.edu/~15849/schedule.html

Fill preference form

- https://forms.gle/JZ93UQvwtepL9KKm7