The Cloud and Virtualization

Agenda

- The cloud
- Who created the cloud
- What is the cloud made
- Virtual machines
  - How far does the rabbit hole go?
  - Case studies

Cloud Computing

- The Cloud
  - A computer utility, warehouse computers; a commodity
  - Catalyst for technology economy
  - Revolutionizing artificial intelligence, machine learning, health care, financial systems, scientific research, and society

Cloud Computing

- The Cloud
  - *ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.*

NIST Cloud Definition

Cloud Computing

- How big is Big Data in the Cloud?
  - Exabytes: Delivery of petabytes of storage daily

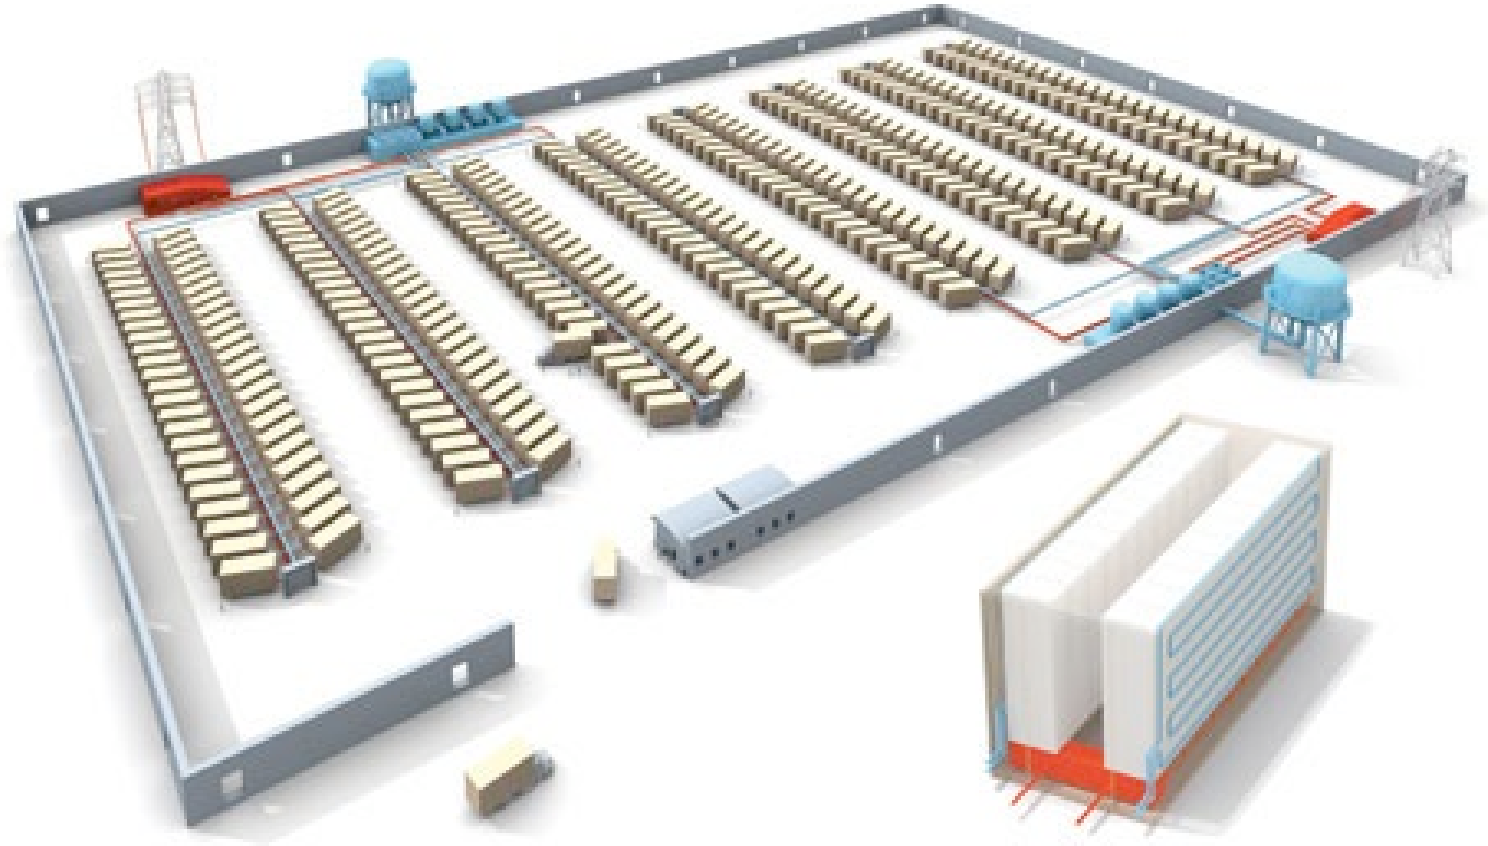

Titan tech boom, randy katz, 2008

Cloud Computing

- How big is Big Data in the Cloud?
  - Most of the worlds data (and computation) hosted by few companies

Cloud Computing

- How big is Big Data in the Cloud?
  - Most of the worlds data (and computation) hosted by few companies

- Currently 4.72 billion internet users
  - 900,000 new users each day [Hootsuite]
- Growing to 175 zettabytes in 2025
- 65% of this data will stored and processed in datacenters [IDC]

Cloud Computing

- How big is Big Data in the Cloud?
  - Most of the worlds data (and computation) hosted by few companies

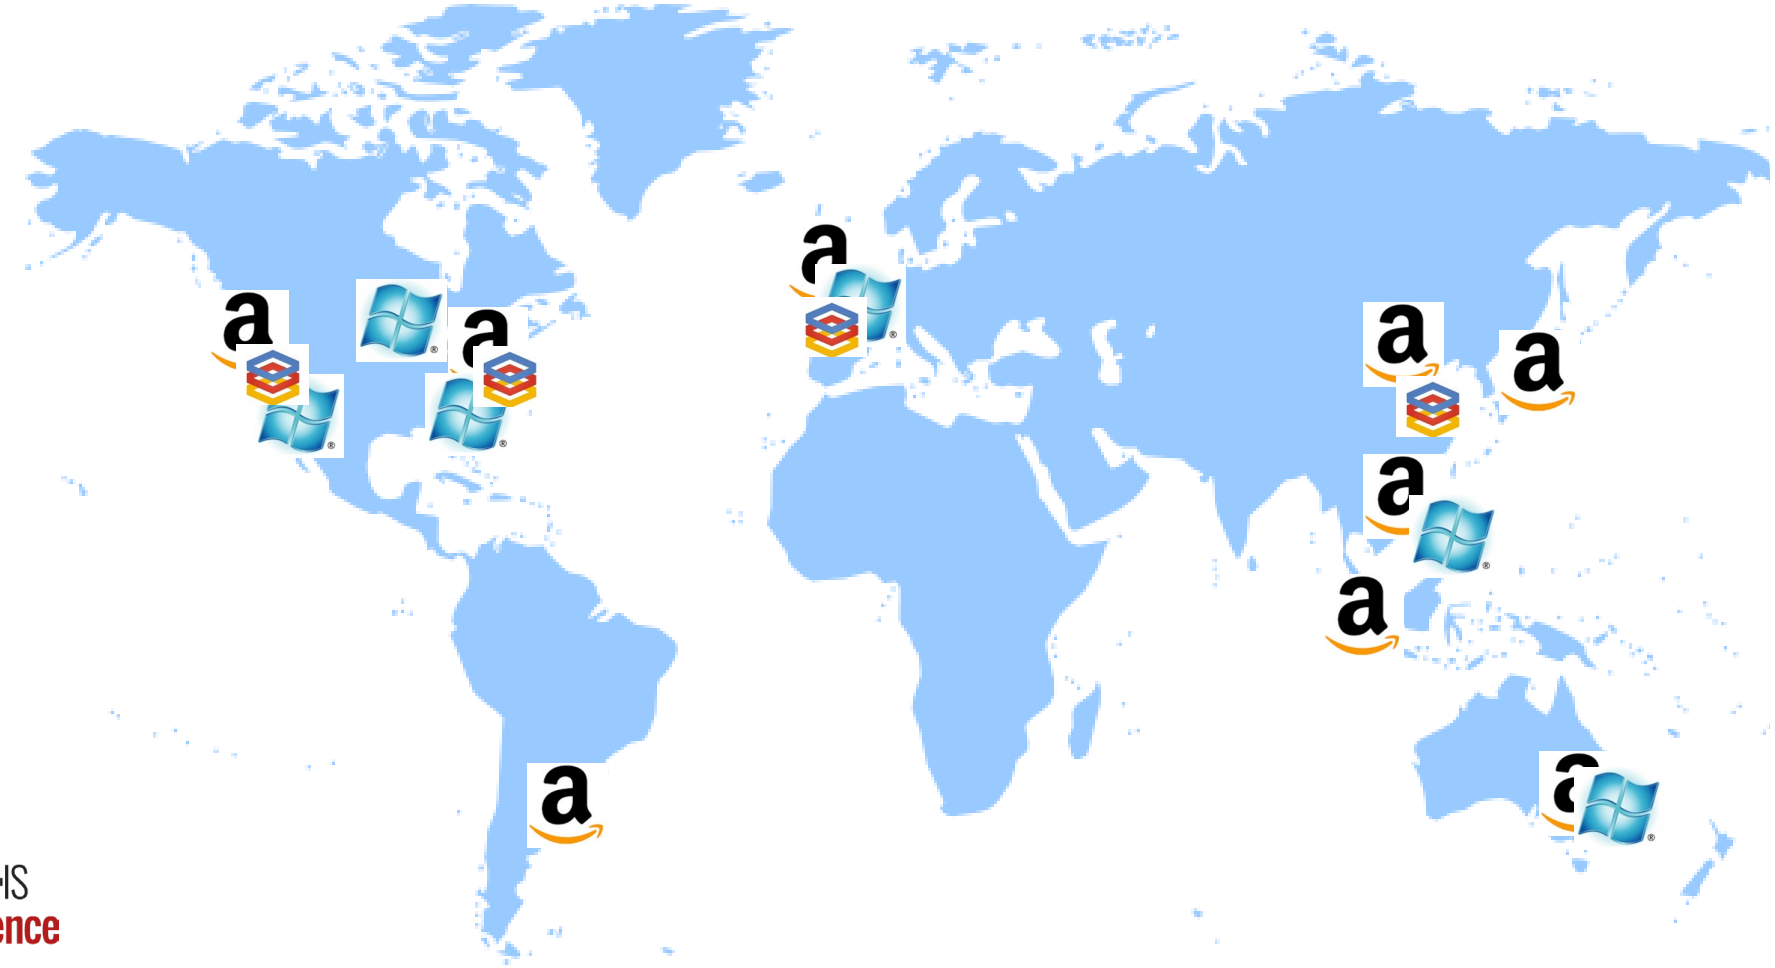

Cloud Computing

- The Cloud
  - A computer utility, warehouse computers; a commodity
  - Catalyst for technology economy
  - Revolutionizing artificial intelligence, machine learning, health care, financial systems, scientific research, and society

Agenda

- The cloud
- Who created the cloud
- What is the cloud made
- Virtual machines
  - How far does the rabbit hole go?
  - Case studies

Who created The Cloud

How Amazon Poached Werner Vogels, 'The Godfather of the Cloud,' From Cornell University

Posted On : January 13, 2023 Published By : Key Executives

Arguably the most famous Chief Technology Officer to ever exist, Werner Vogels has spent the past 17 years transforming Amazon from an online book seller into one of the largest ecommerce empires on the globe. A native of the Netherlands, Vogels is celebrated as a top expert on ultra-scalable systems. He holds a doctorate in computer science and is the author of many conference and journal articles, most of them on computer-related distributed systems.

Search ... Search

CATEGORIES

Select Category

RECENT POSTS

Amy Reichanadter Shapes Databricks' People-Centric Growth Strategy May 1, 2025

Building More Than Bricks: LEGO's New Factory Powers Sustainability and Supports Education in Vietnam May 1, 2025

André Heinz's Global Perspective on Leading People Through Change at Celonis April 30, 2025

Building Plaid: Zach Perret's Mission to Create an Open, Inclusive Financial System April 30, 2025

Inside Bhawna Singh's Journey from Engineer to Global CTO Leader April 29, 2025

Dinni Jain's Journey from Wall Street to the Fields and Fiber Lines of America April 28, 2025

Who created The Cloud

- Robbert van Renesse, Ken Birman, Werner Vogels
  - Researchers at Cornell sold company to Amazon

- Werner Vogels
  - Turns Amazon Web Services (AWS) into the cloud
  - Idea: Sell excess computing during after Christmas season

Agenda

- The cloud
- Who created the cloud
- What is the cloud made
- Virtual machines
  - How far does the rabbit hole go?
  - Case studies

Big Picture: Recall (Virtual) Memory

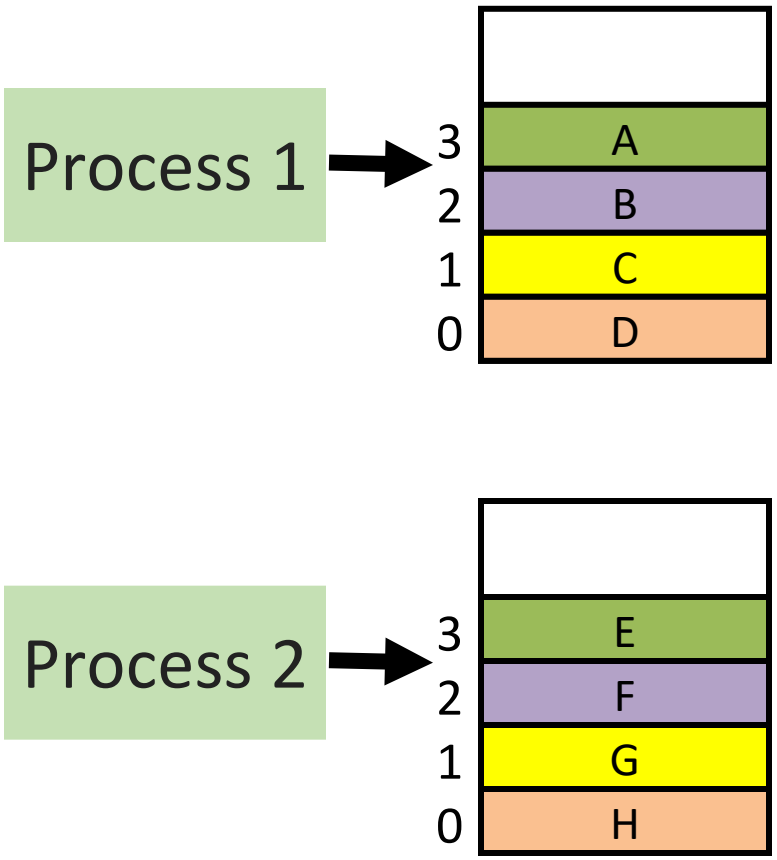

Give each process an illusion that it has exclusive access to entire main memory

Big Picture: Recall (Virtual) Memory

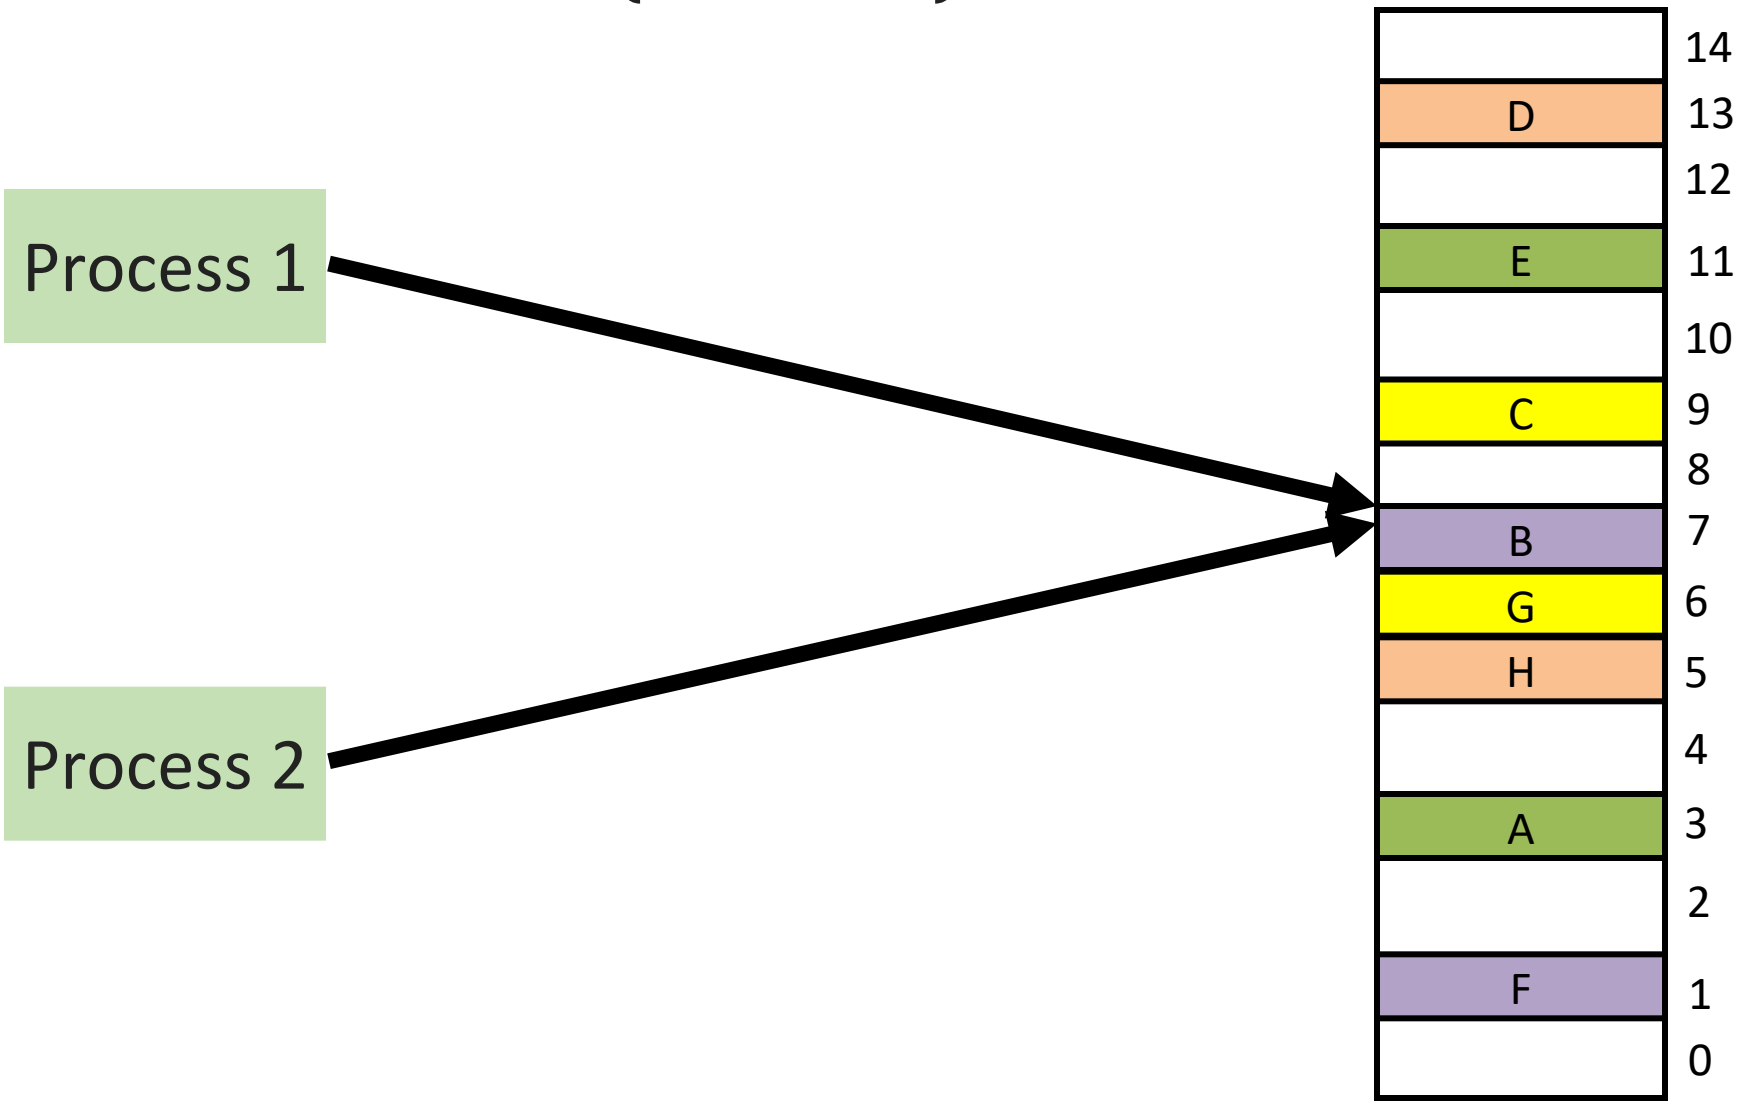

How do we create the illusion?

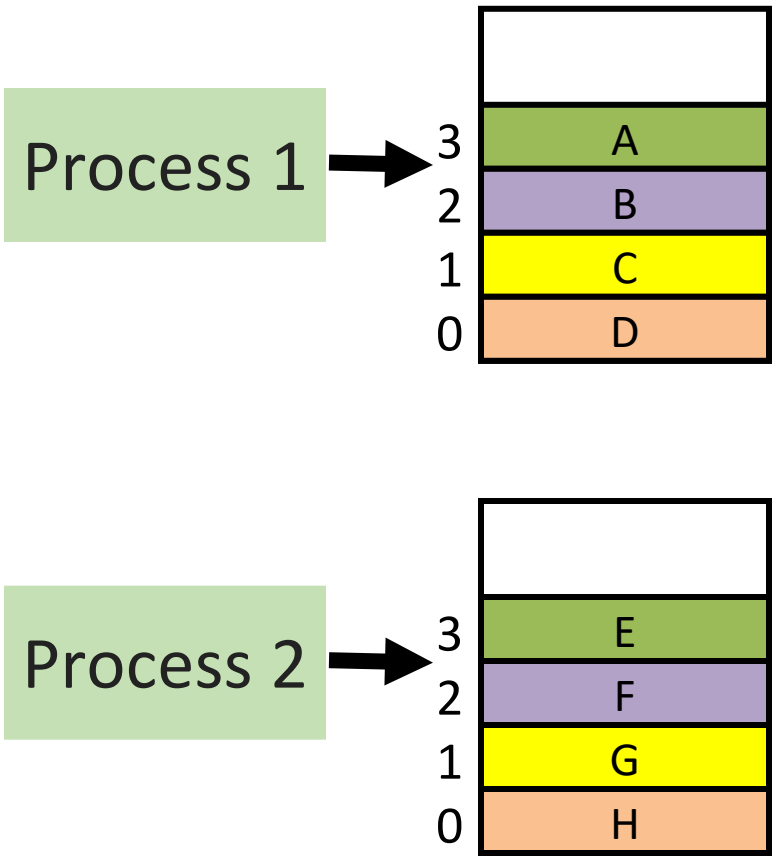

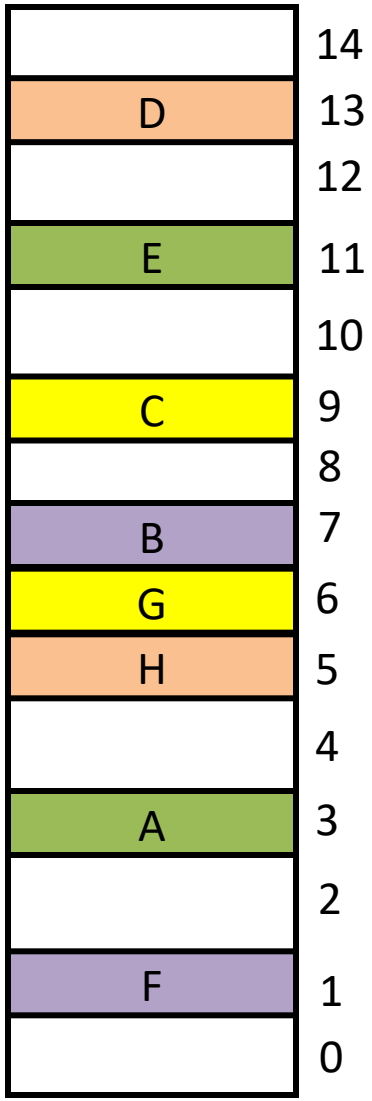

Physical Memory

How do we create the illusion?

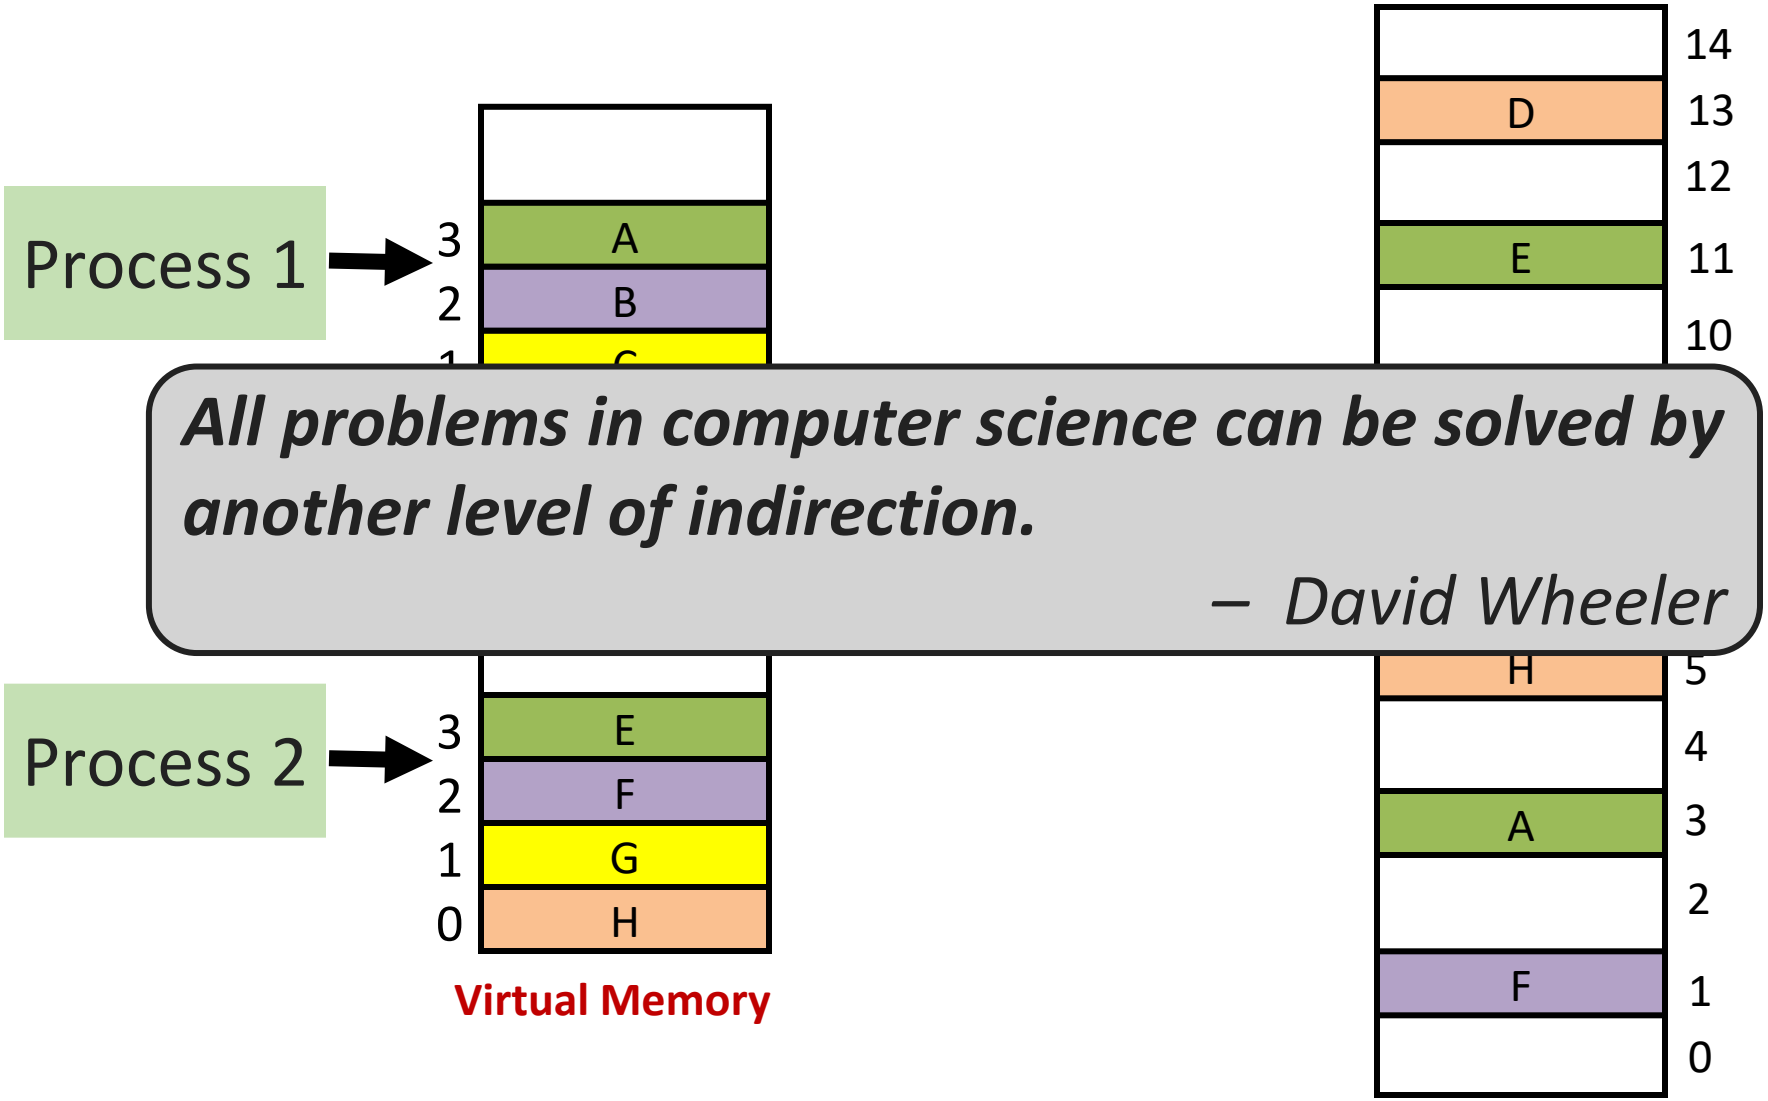

All problems in computer science can be solved by another level of indirection.

– David Wheeler

How do we create the illusion?

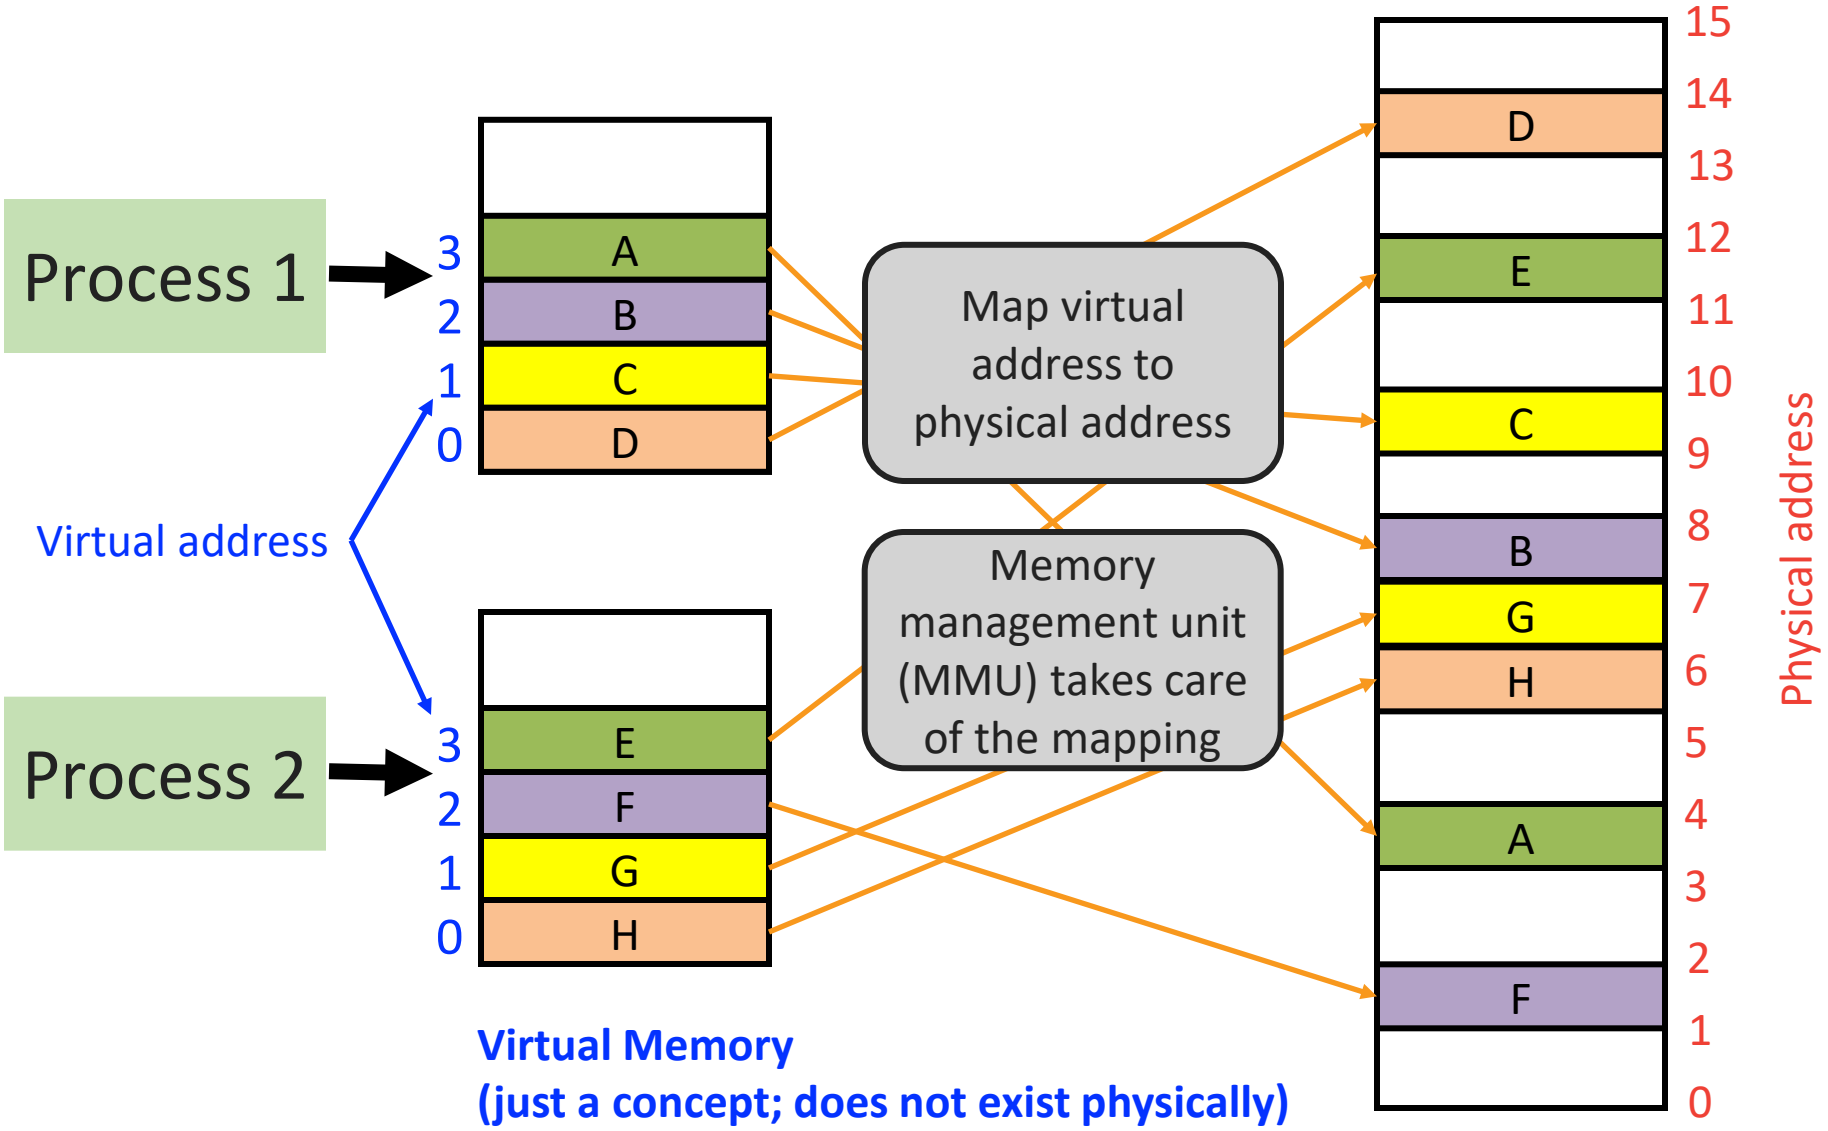

What if physical memory was virtual?!

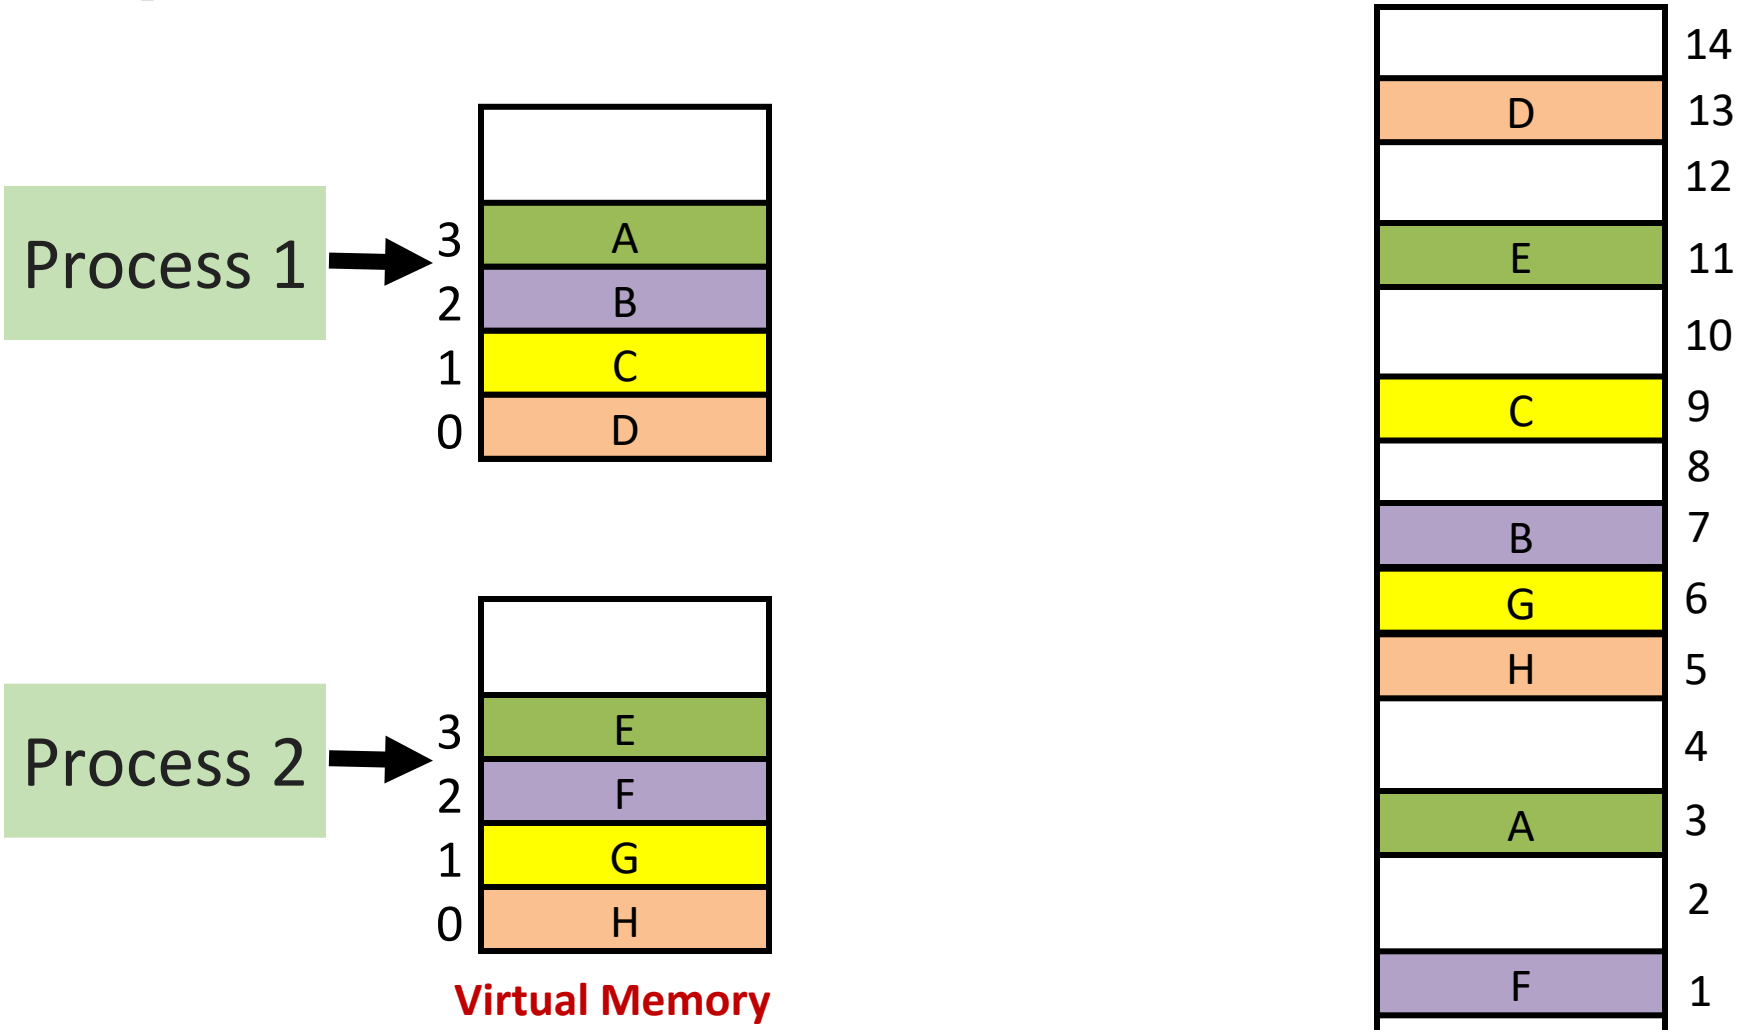

What if physical memory was virtual?!

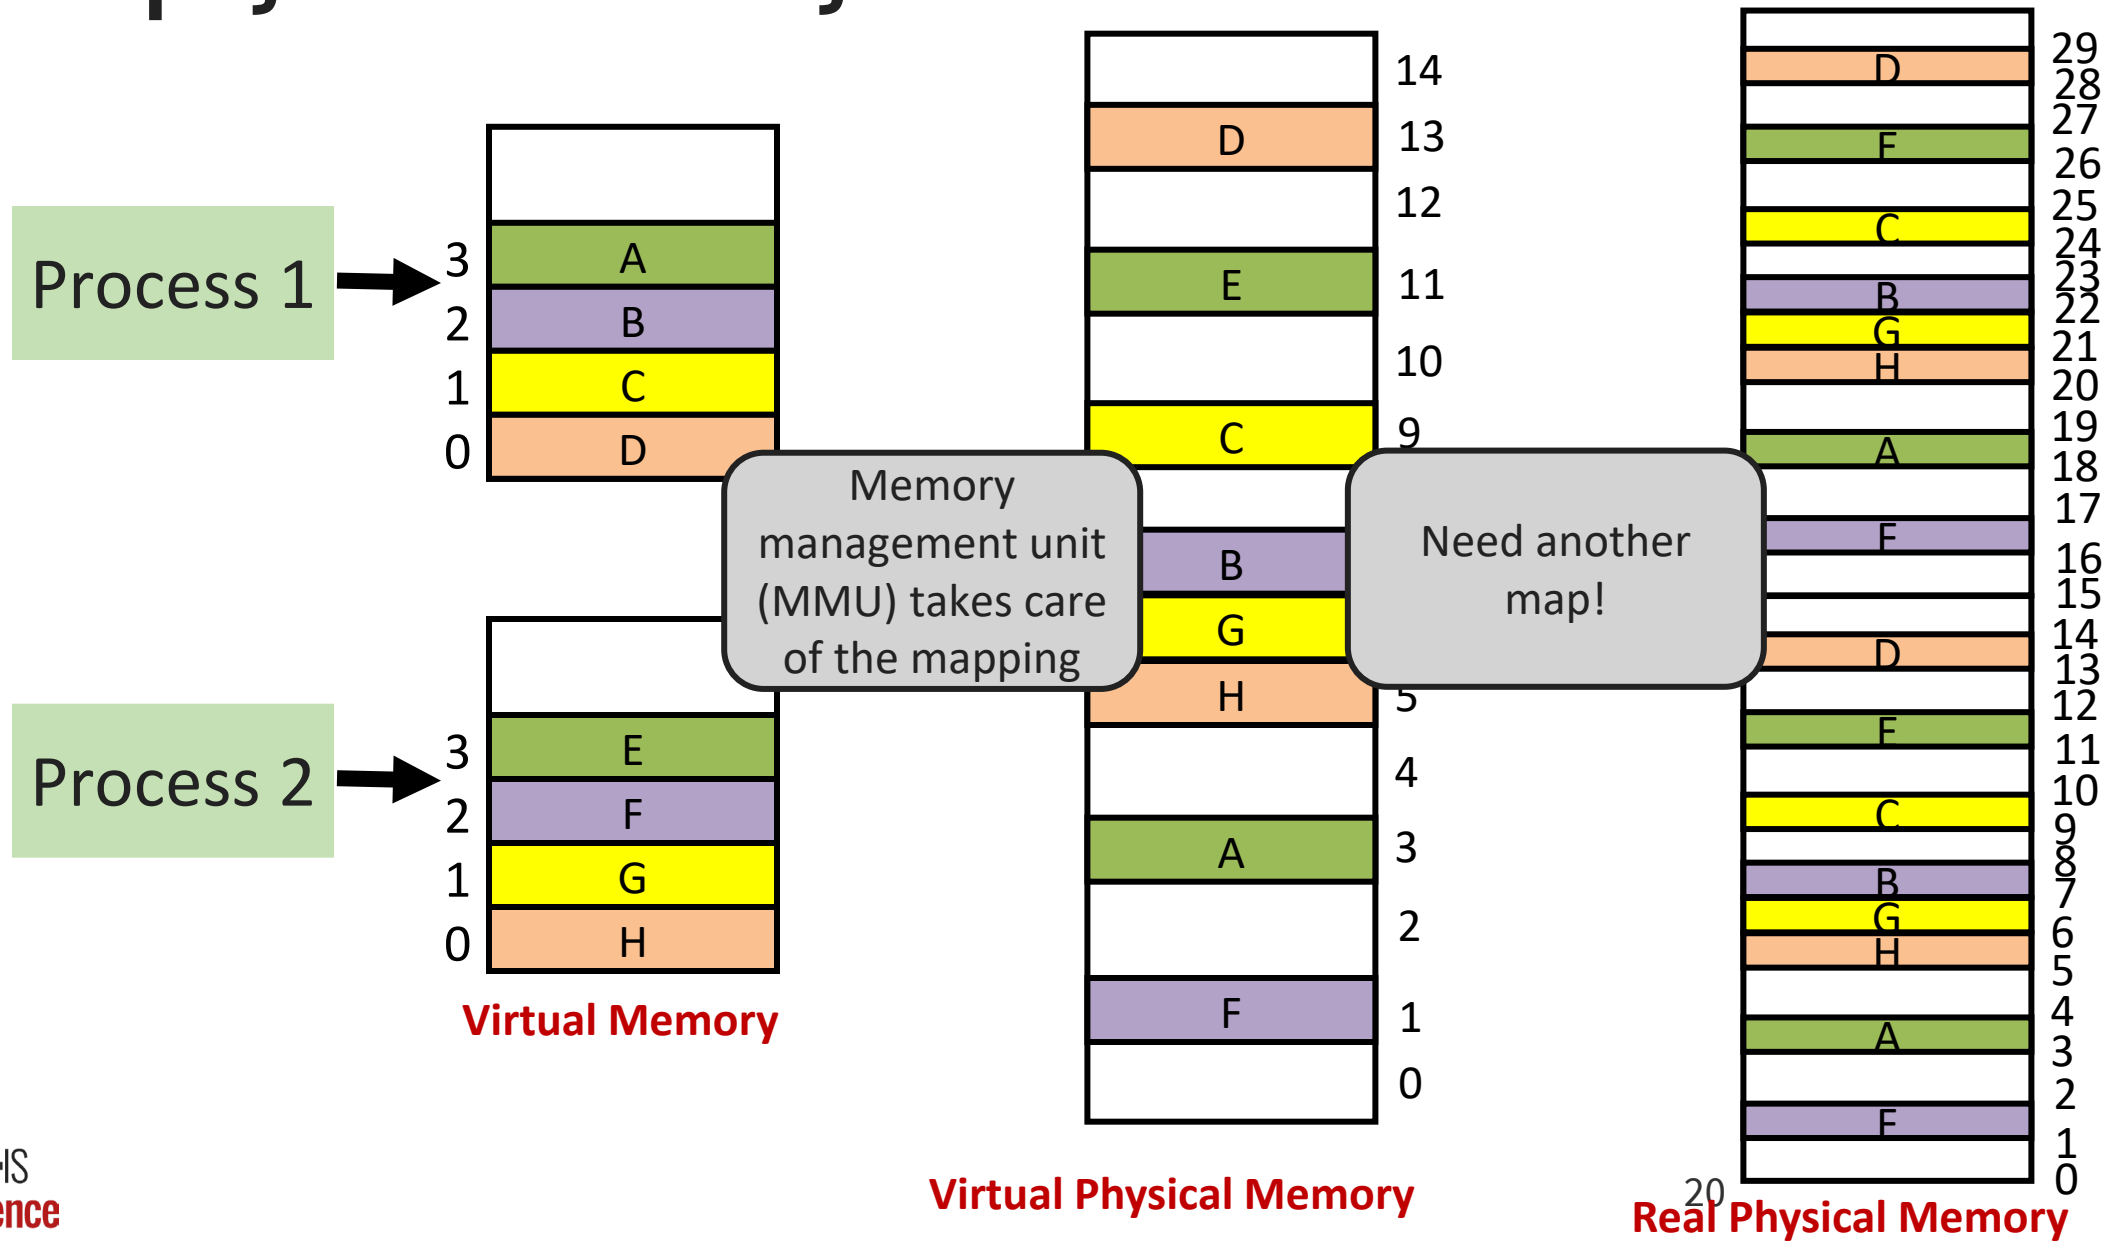

The Cloud and Virtualization

- Operating System Kernel
  - Manages hardware resources and (virtual to physical) memory
- Hypervisor
  - Manages hardware resources and (virtual to physical) memory
  - OS kernel memory is virtual!
  - In general, a virtual machine (VM) is the emulation of a computer system

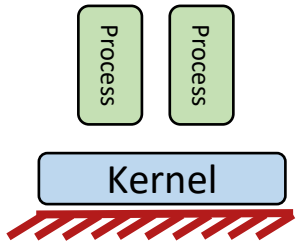

Processes + Operating System + virtual memory

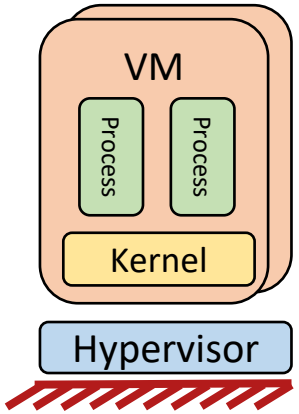

Virtual Machine: Hypervisor + Processes + Operating System + virtual memory

The Cloud and Virtualization

- Operating System Kernel
  - Manages hardware resources and (virtual to physical) memory
- Hypervisor (a.k.a. virtual machine monitor [VMM])
  - Manages hardware resources and (virtual to physical) memory
  - OS kernel memory is virtual!
  - In general, a virtual machine (VM) is the emulation of a computer system
- Limits of virtualization
  - Practically, 2 or 3 level of virtualization
  - Theoretically, unlimited
    - Turtles/IBM

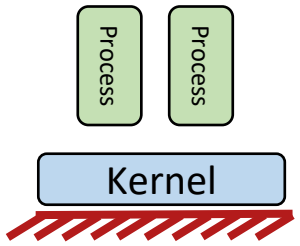

Processes + Operating System + virtual memory

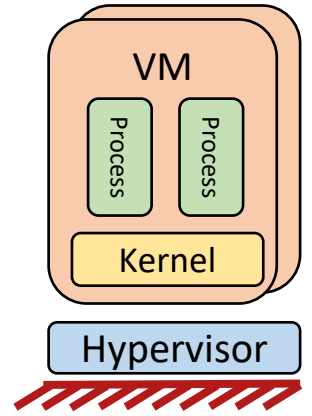

Virtual Machine: Hypervisor + Processes + Operating System + virtual memory

The Cloud and Virtualization

- Operating System Kernel
  - Manages hardware resources and (virtual to physical) memory
- Hypervisor (a.k.a. virtual machine monitor [VMM])
  - Manages hardware resources and (virtual to physical) memory
  - OS kernel memory is virtual!
  - In general, a virtual machine (VM) is the emulation of a computer system
- Limits of virtualization
  - Practically, 2 or 3 level of virtualization
  - Theoretically, unlimited
    - Turtles/IBM
- Containers (e.g. Docker)
  - Provide portability

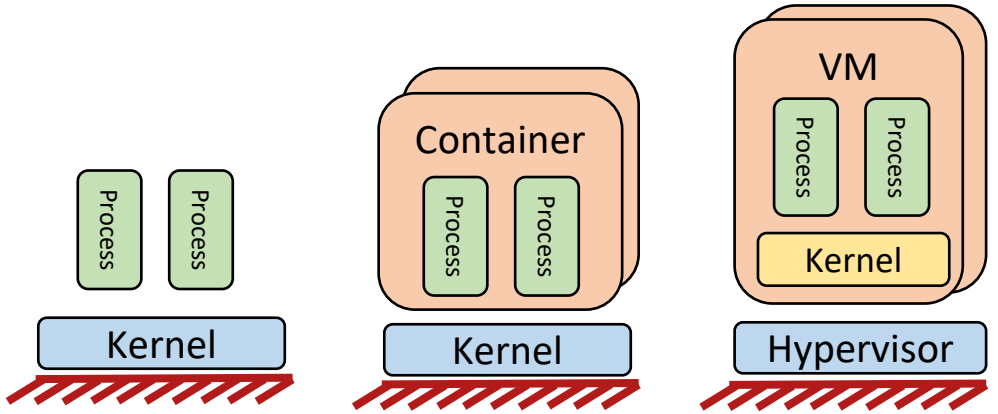

Processes + Operating System + virtual memory

Container

Virtual Machine: Hypervisor + Processes + Operating System + virtual memory

Who created The Cloud: Revisited

- Robbert van Renesse, Ken Birman, Werner Vogels
  - Researchers at Cornell sold company to Amazon

- Werner Vogels
  - Turns Amazon Web Services (AWS) into the cloud
  - Idea: Sell excess computing during after Christmas season

  - Werner downloaded a virtual machine (VM) hypervisor, Xen
  - Xen was used to partition a physical machine into many virtual machines (VMs)
  - AWS sells VMs!
  - This is the cloud!

Agenda

- The cloud
- Who created the cloud
- What is the cloud made
- Virtual machines
  - How far does the rabbit hole go?
  - Case studies

Virtualization

- A virtual machine (VM) is the emulation of a computer system

- Why virtualize
  - Underutilized machines
  - Easier to debug and monitor operating system (OS)
  - Portability
  - Isolation
  - The cloud (e.g. Amazon EC2, Google Compute Engine, Microsoft Azure)

Virtualization

- 1960's: first track of virtualization
  - Time and resource sharing on expensive mainframes
  - IBM VM/370
- Late 1970's and early 1980's: became unpopular
  - Cheap hardware and multiprocessing OS
- Late 1990's: became popular again
  - Wide variety of OS and hardware configurations
  - VMWare
- Since 2000: hot and important
  - Cloud computing
  - Docker containers

Full Virtualization

- Complete simulation of underlying hardware
- Unmodified guest OS
- Trap and simulate privileged instruction
- Guest OS can't see real resources

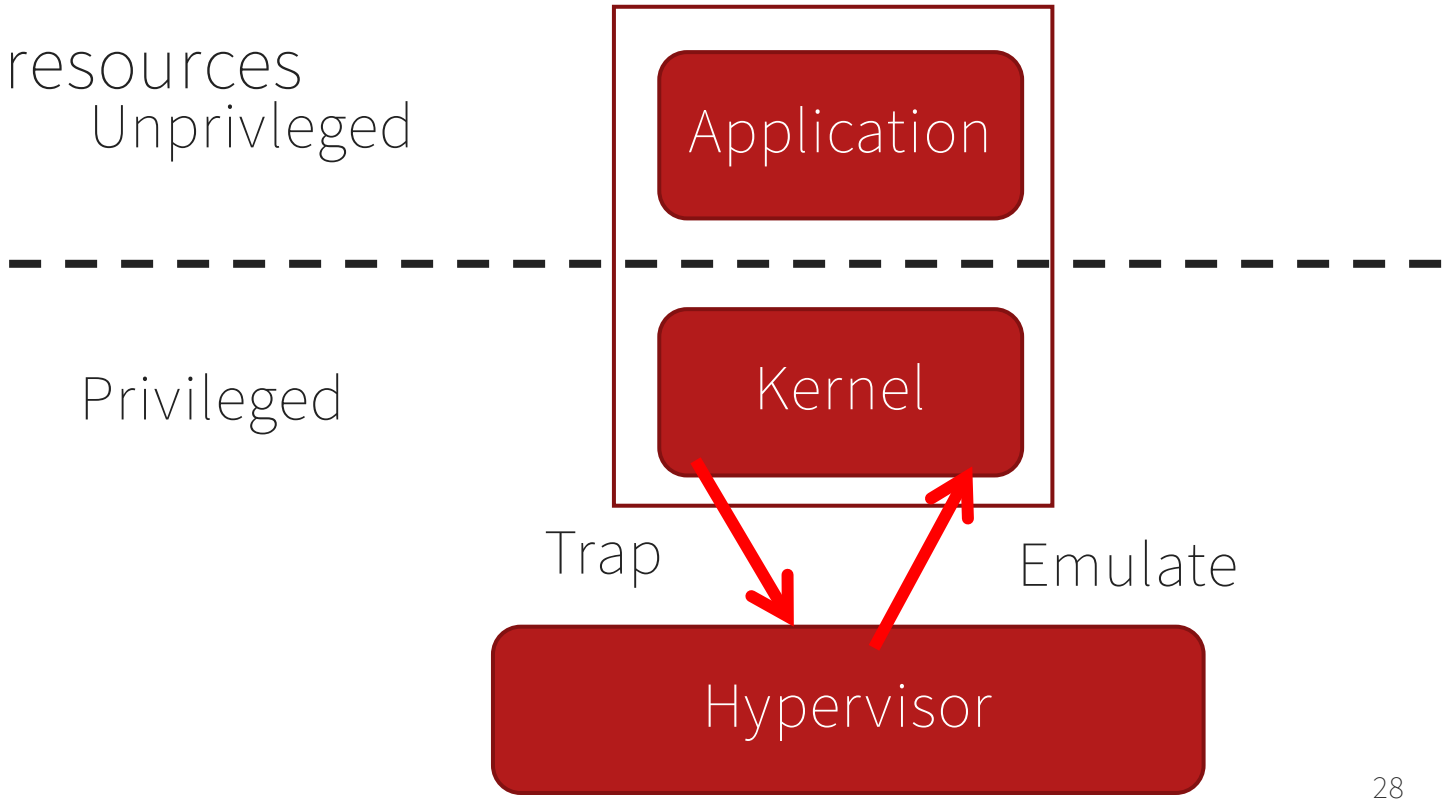

Cool Virtualization Capabilities

- Operating System can freeze (context switch) a process
  - Process control block (PCB) describes memory and state of processor
- A hypervisor can freeze a virtual machine (VM)!
  - Hypervisor controls state of memory and processor
- A hypervisor can move/migrate a VM via Virtualization
  - Virtualize processor Instruction Set Architecture (ISA) and memory
  - Similar to paging, migrate memory between machines

Agenda

- The cloud
- Who created the cloud
- What is the cloud made
- Virtual machines
  - How far does the rabbit hole go?
  - Case studies

Case Study: Xen

- Xen
  - DomU hosts guest operating system in virtual machine
  - Dom0 manages devices and guests
  - Control Transfer: Hypercalls and Events

    (like syscalls and device interrupts)

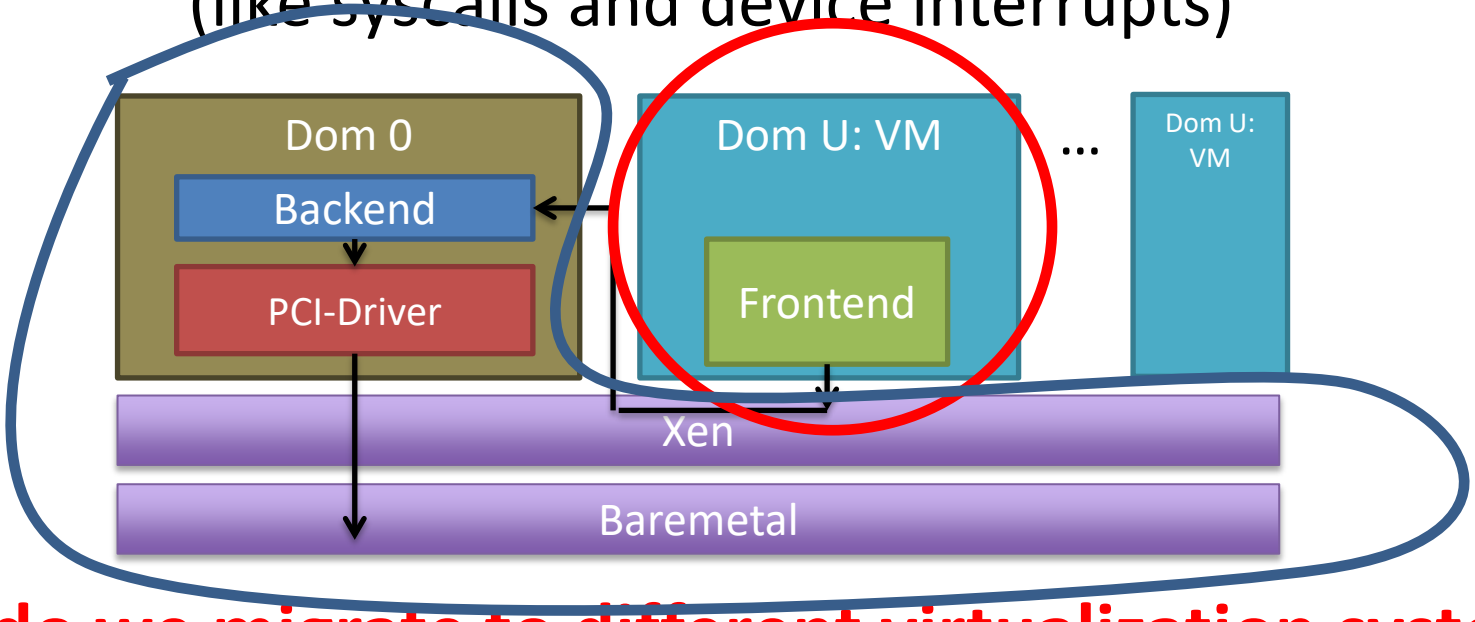

How do we migrate to different virtualization system?

Case Study: Xen

- Virtualization enables migration of computation across a single cloud

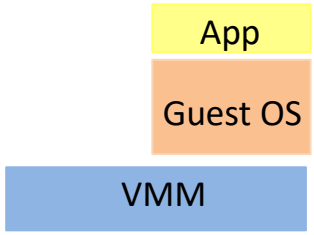

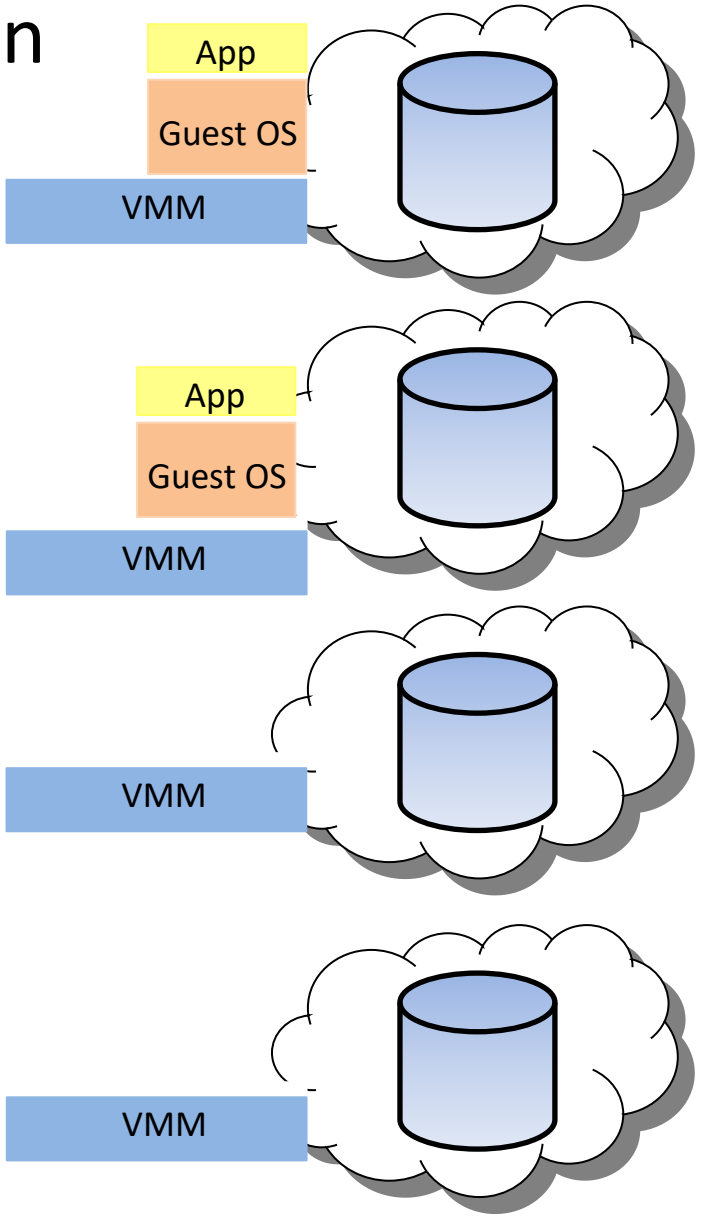

Case Study: Xen

- What about migrating across the cloud?

Case Study: Xen-Blanket

- Need another layer of virtualization (indirection)

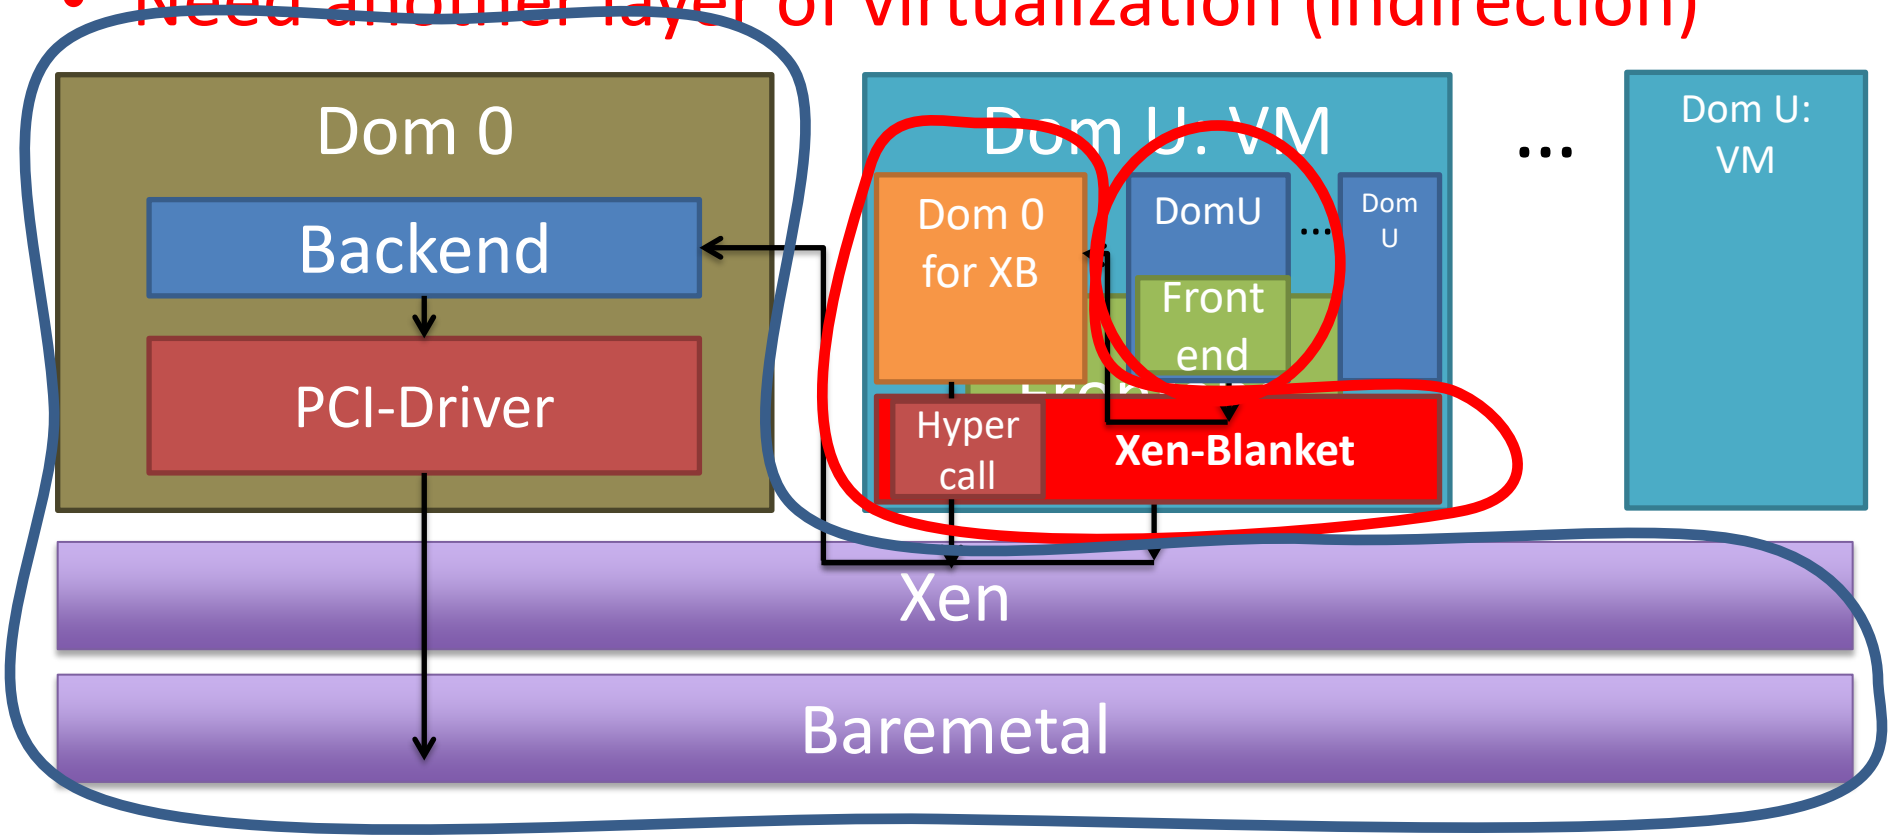

- Blanket drivers enable access to baremetal
- Blanket drivers support memory address translation
- Hypercalls provides privilege support for HVM

Case Study: Xen-Blanket

- Can create your own
  Cloud-within-a-Cloud
  aka a "Supercloud"

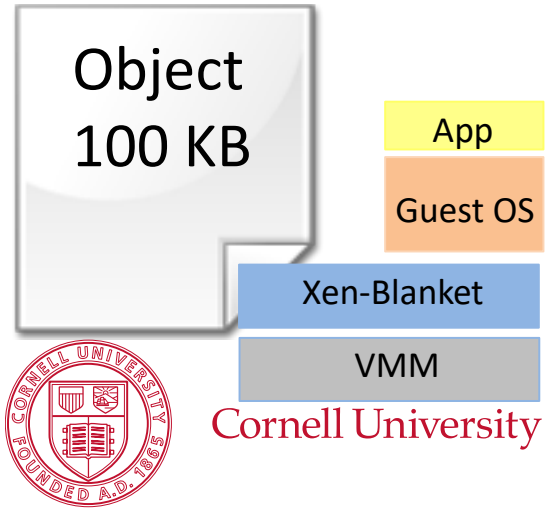

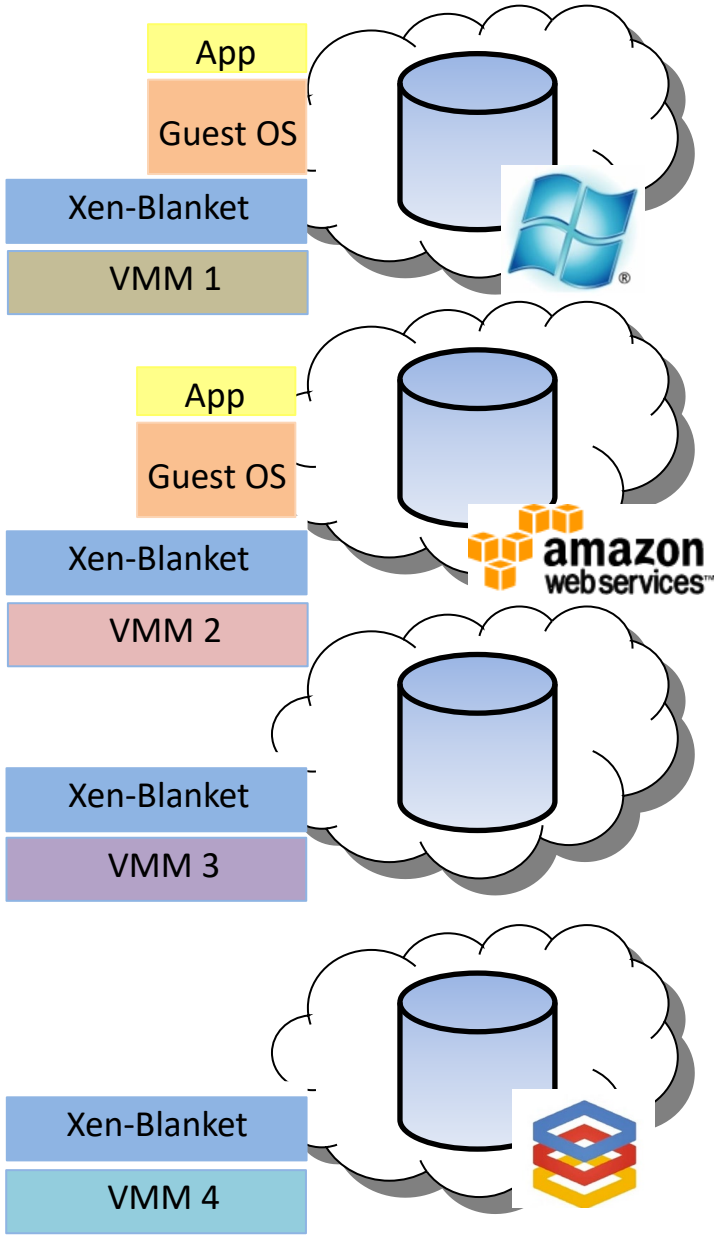

- Migrate computation among different cloud providers

Case Study: Xen-Blanket

- Migration of computation among different cloud providers
- Small overhead
- Migration: 1 second downtime

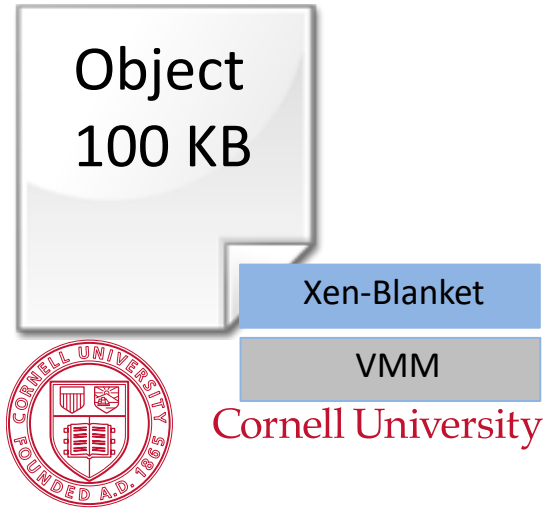

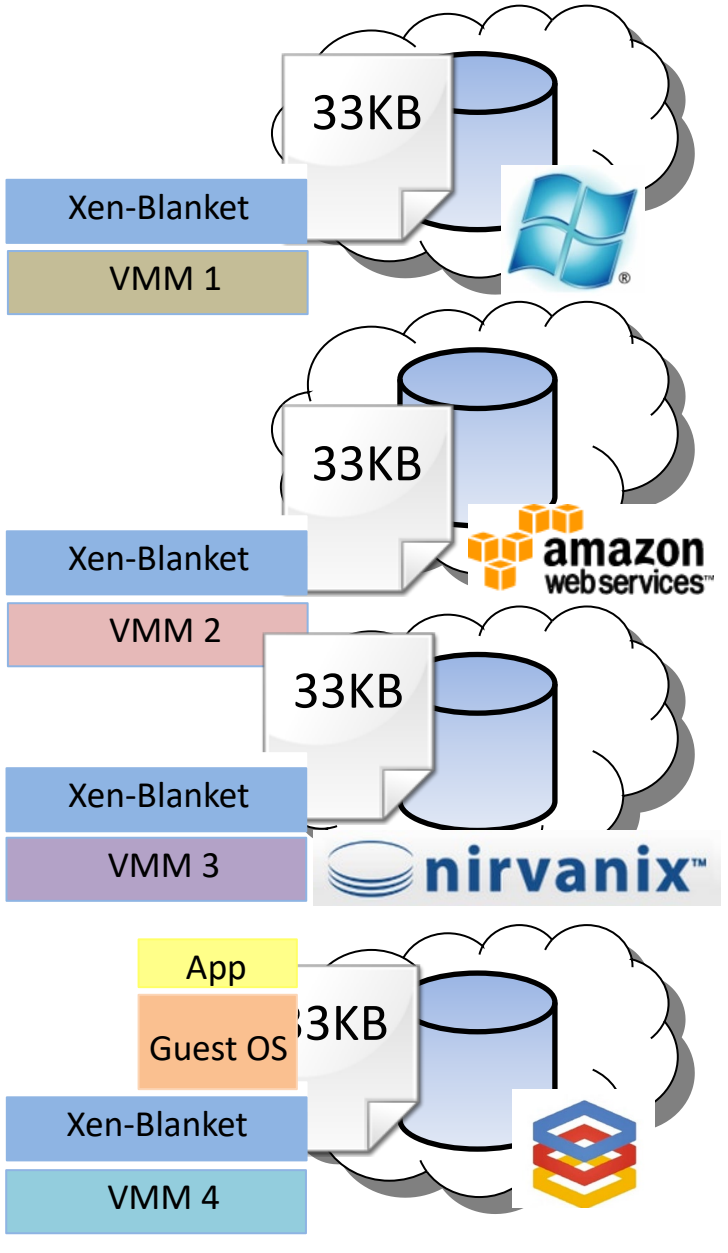

- ***Is VM migration sufficient?***

Case Study: Xen-Blanket

- How to build a "Supercloud" and "virtual wires"?

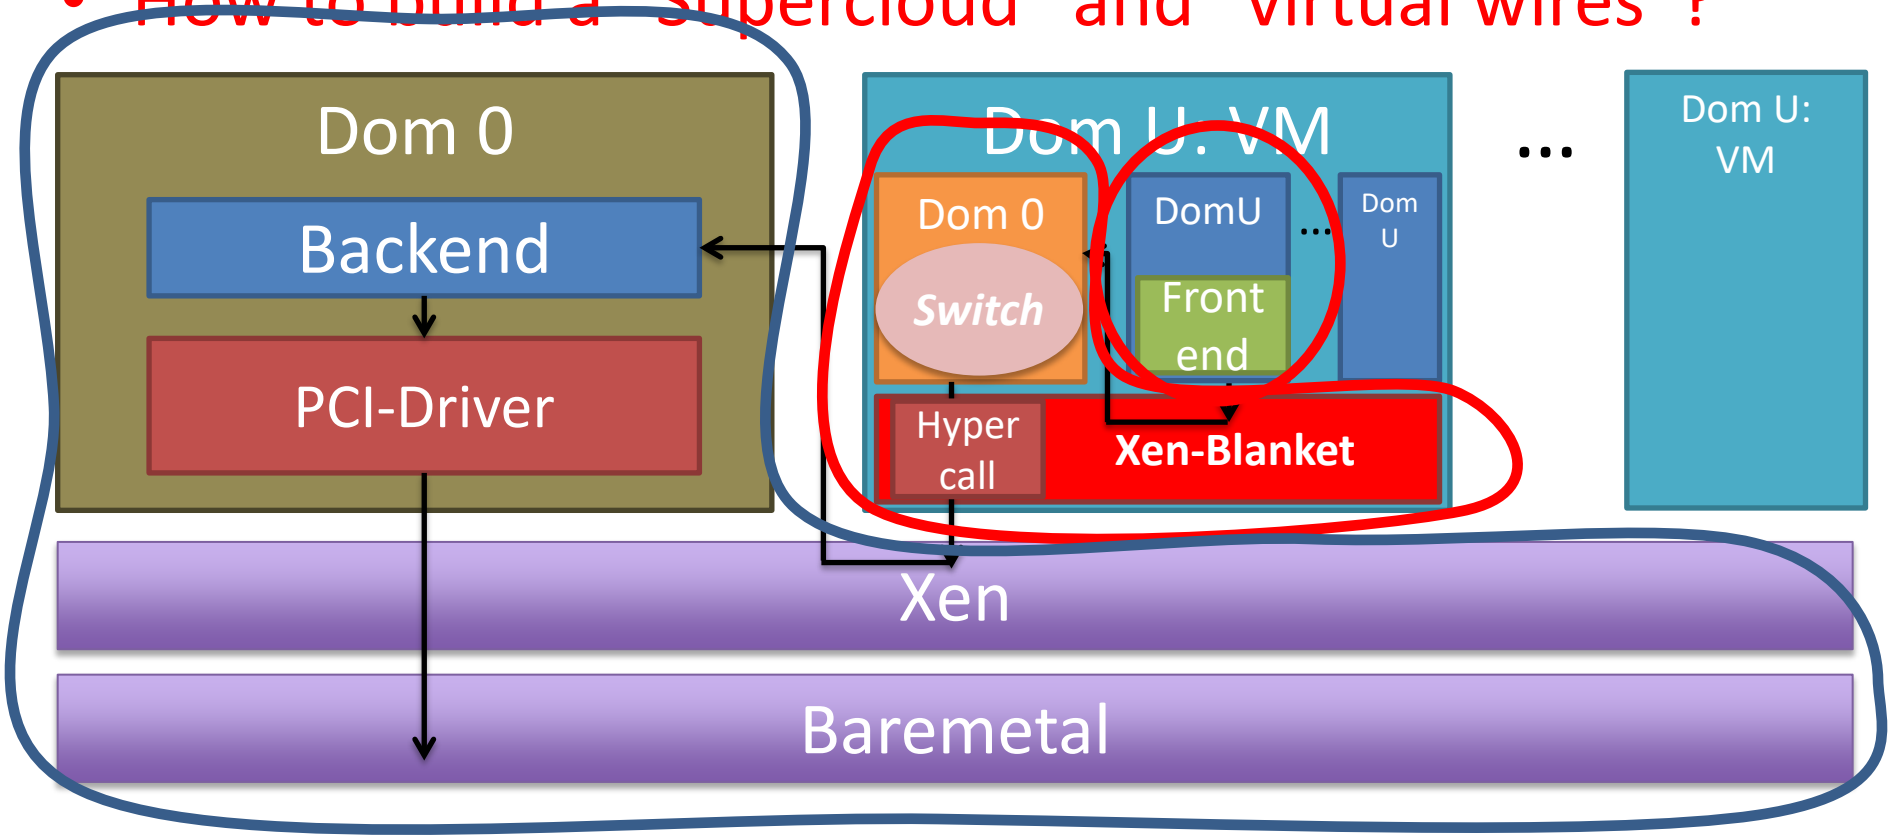

- Blanket drivers enable access to baremetal
- Blanket drivers support memory address translation
- Hypercalls provides privilege support for HVM

Case Study: Xen-Blanket

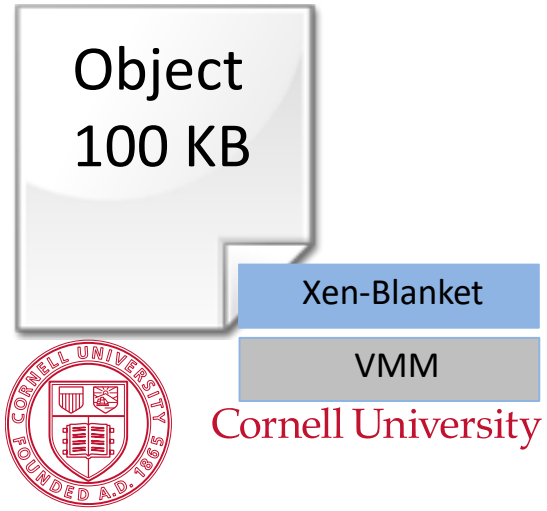

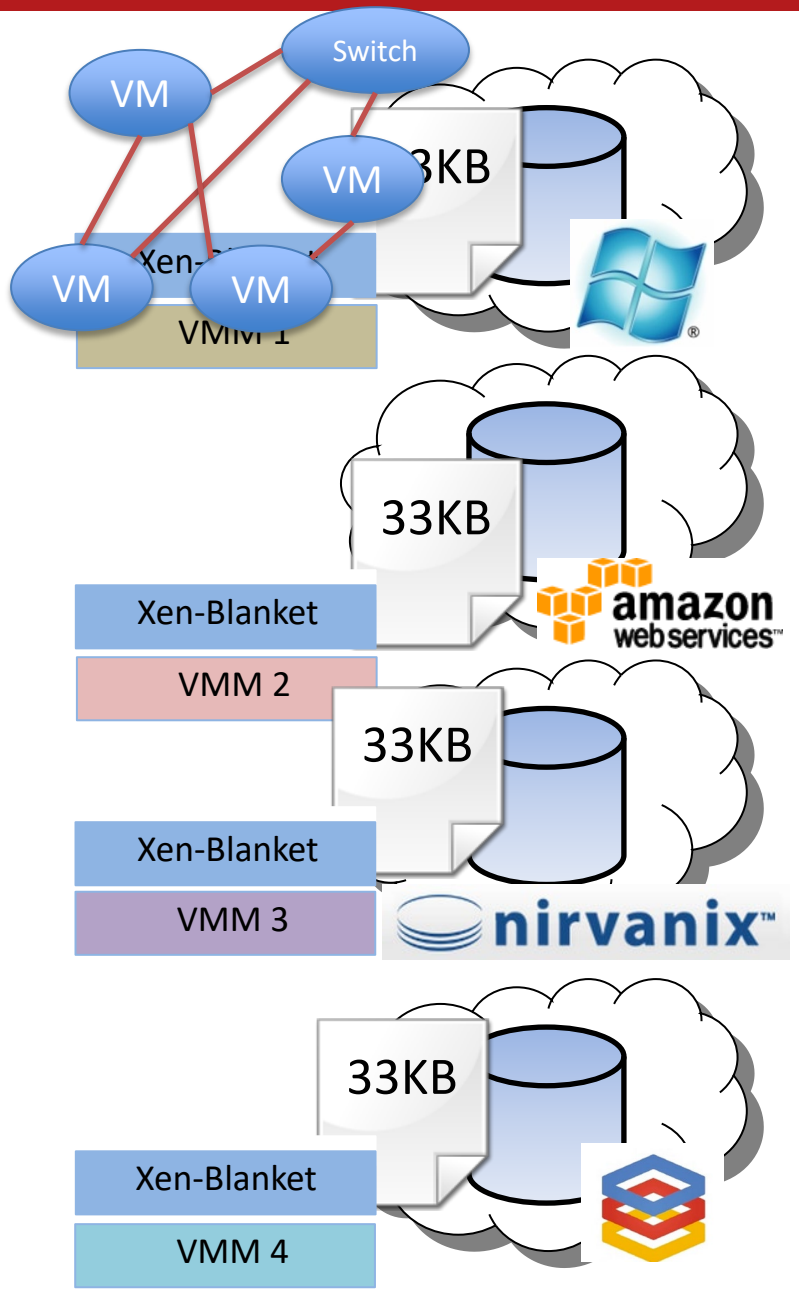

Case Study: Xen-Blanket

- Virtual Wire
  - Uses virtual switches
  - Installation of virtual connectors
  - Encapsulates and sends packets

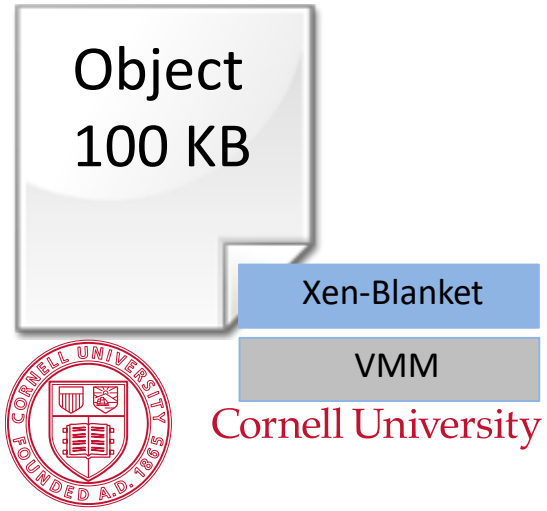

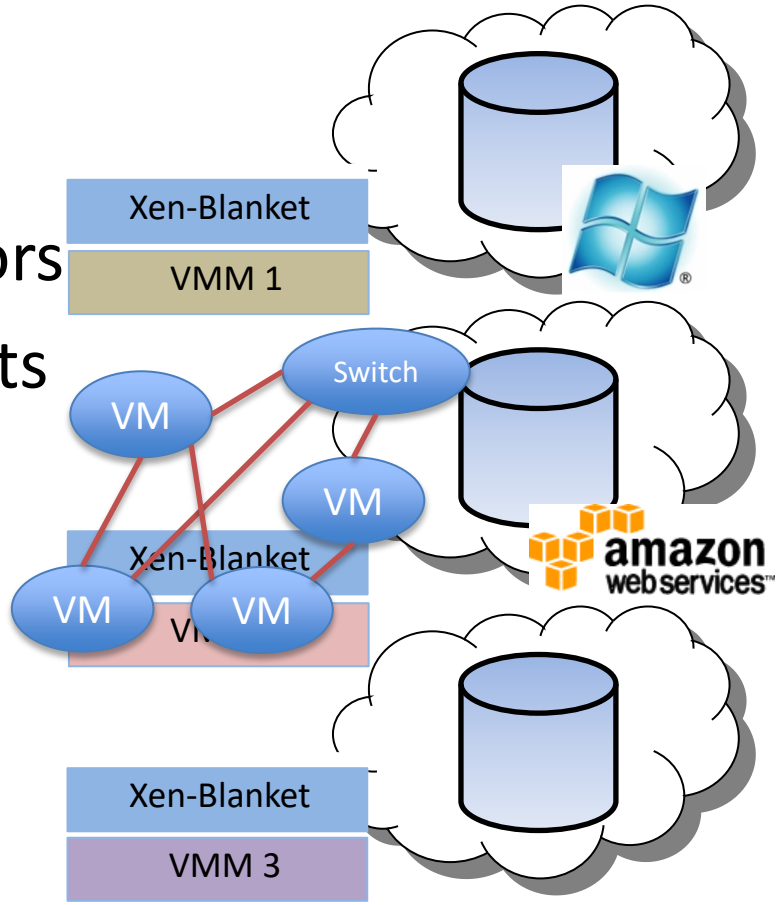

- Migration of Network Configs
  - Migrates network topology
  - Obviates administrative configs

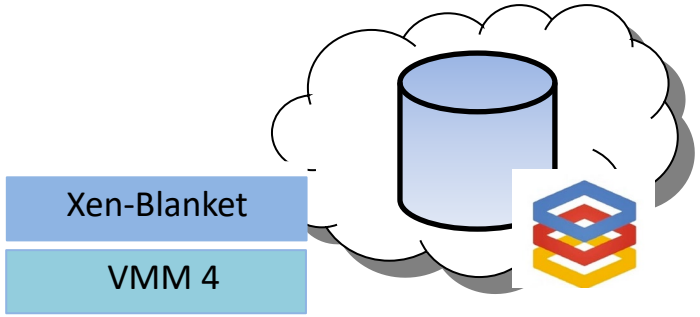

Case Study: Spot Market

- Often very cheap, sometimes 10x cheaper!
- Price changes dynamically

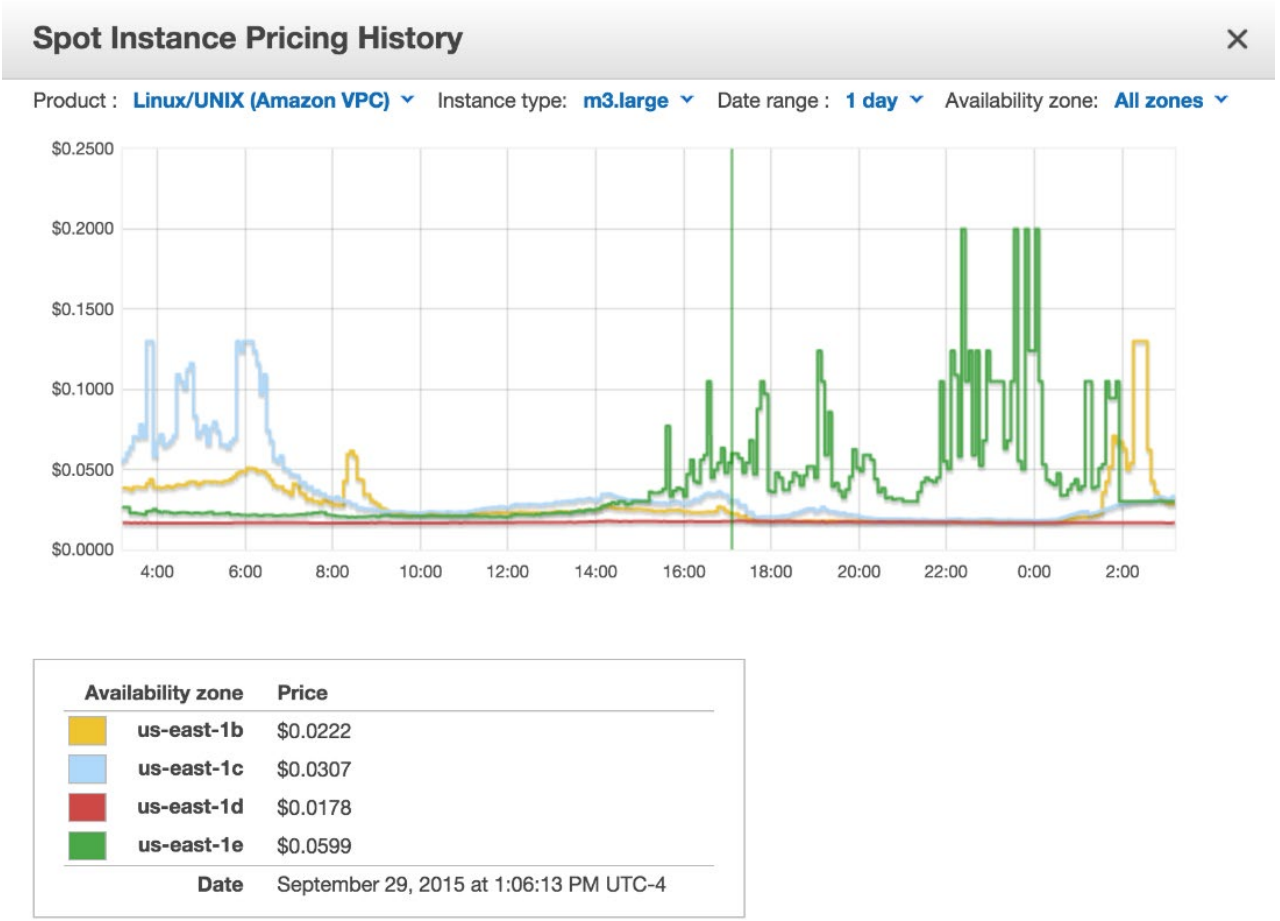

Case Study: Spot Market

- Often very cheap, sometimes 10x cheaper!
- Price changes dynamically
- Problem: Unreliable
  - Applications terminated with little to no warning
  - Applications need to handle failures
  - Very hard to adopt for legacy and stateful applications

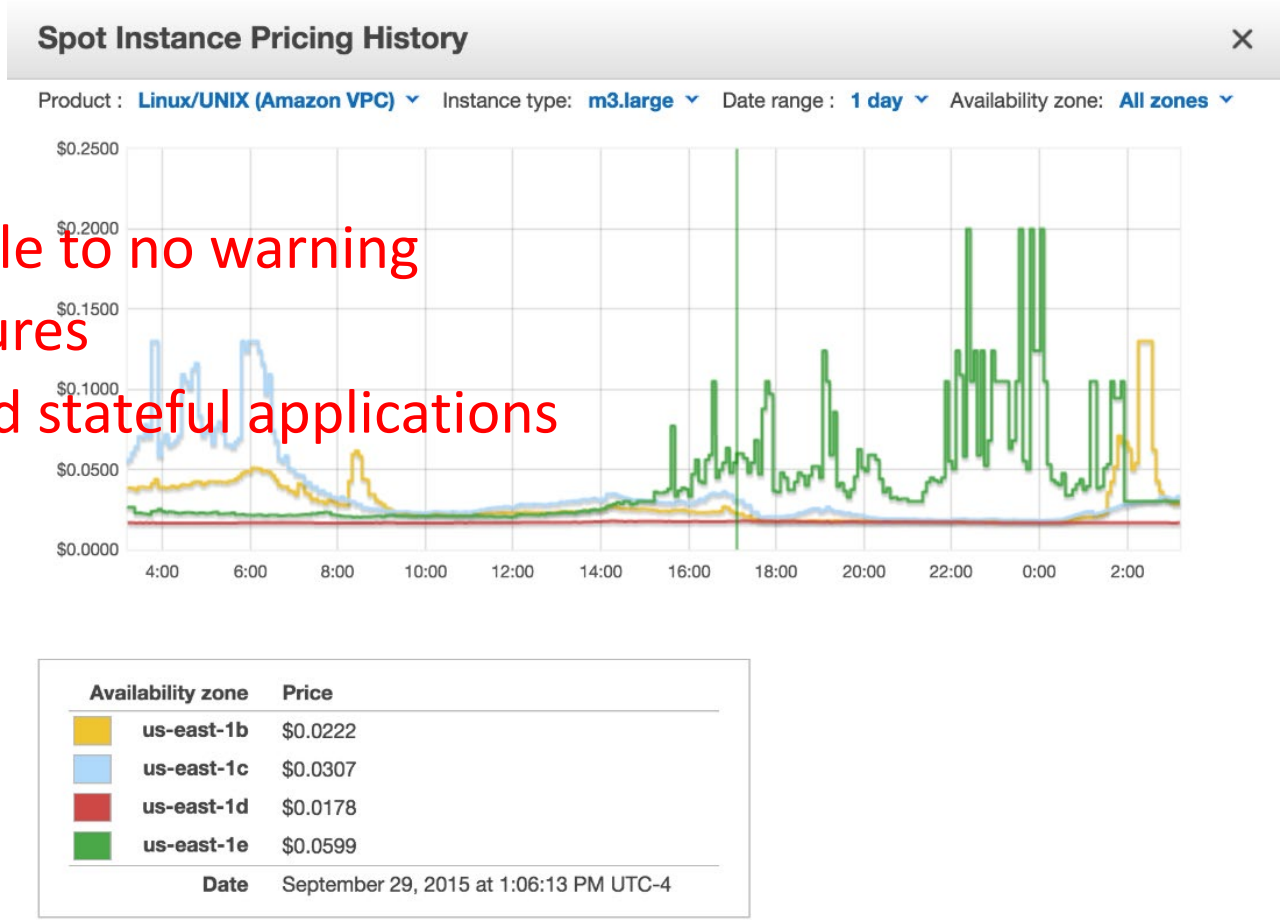

Case Study: Spot Market

Use virtualization to

- Reduce the cost of using the cloud via the Spot market
- Improve the reliability of (stateful) applications
- Don't require *any* changes to the application
- Achieve good performance

Approach

- Live virtual machine migration of (stateful) applications
- Problem: Cloud user cannot live migrate applications
  - Cloud providers own the physical machine, not Cloud users

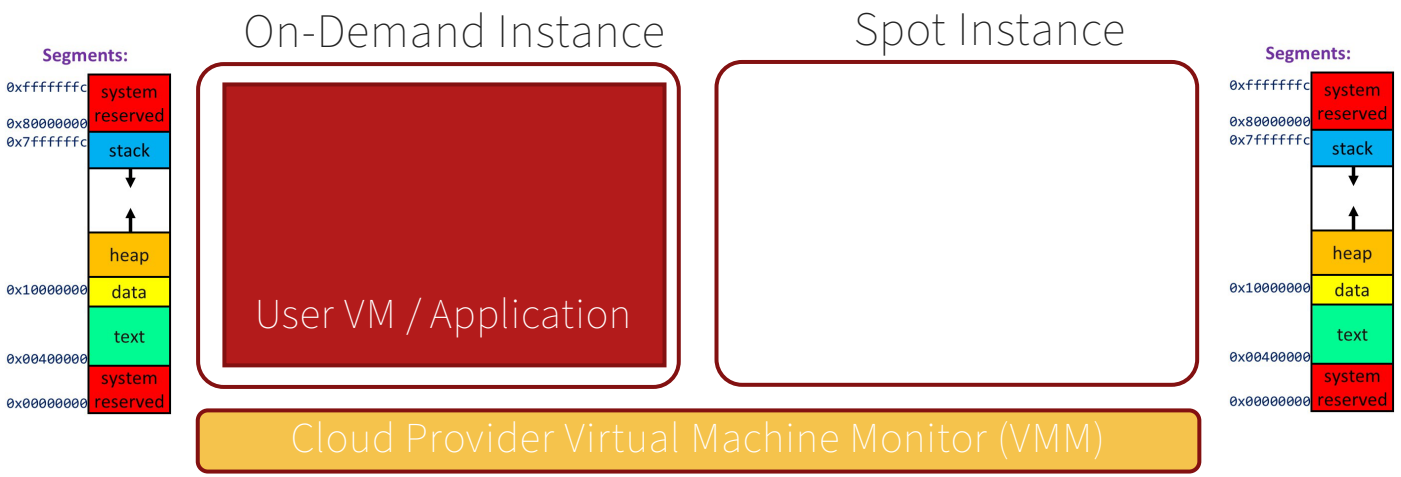

Approach

- Live virtual machine migration of (stateful) applications
- Problem: Cloud user cannot live migrate applications
  - Cloud providers own the physical machine, not Cloud users

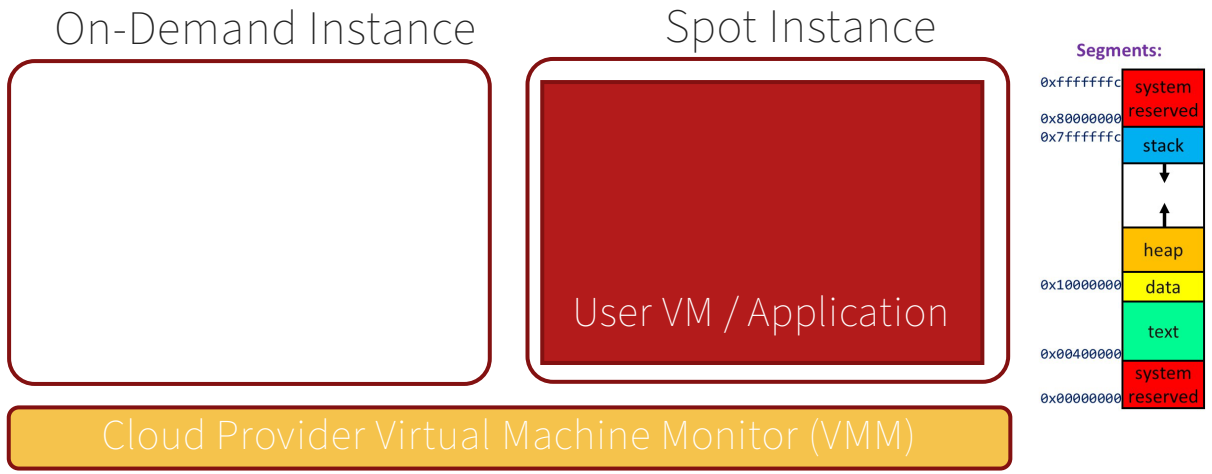

Solution

- Nested Virtualization

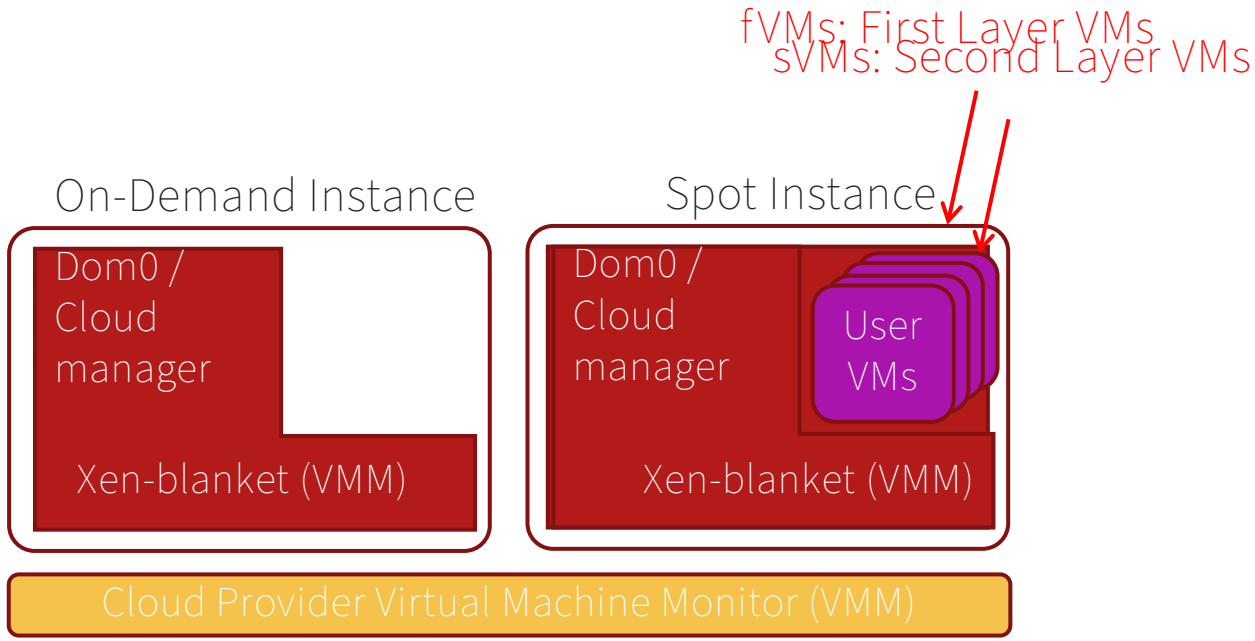

See *The Xen-Blanket: Virtualize Once, Run Everywhere*
D. Williams, H. Jamjoom, and H. Weatherspoon. In ACM Eurosys, April 2012

Solution

- Nested Virtualization
- Can even migrate across availability zones

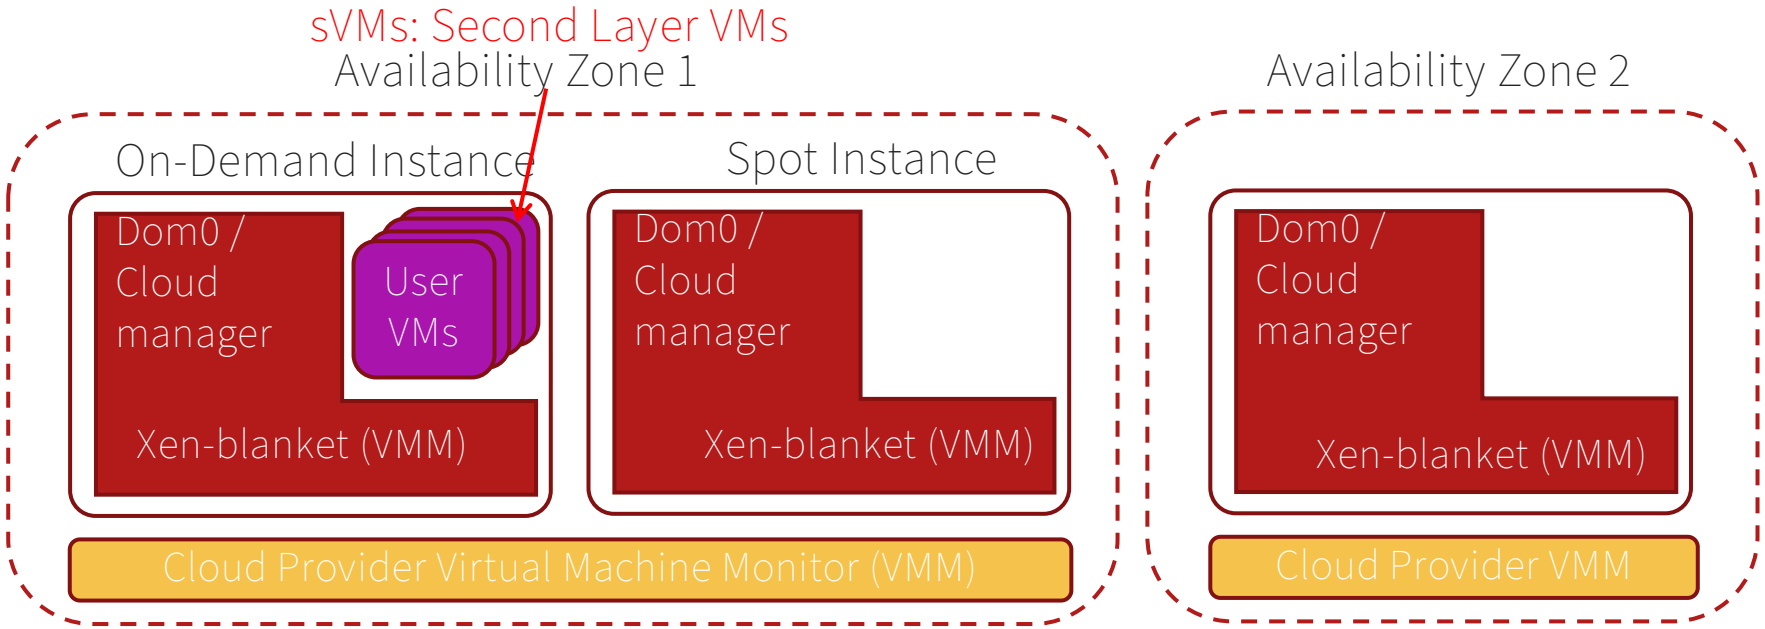

See *Supercloud: A Library Cloud for Exploiting Cloud Diversity.* Z. Shen, Q. Jia, G.E. Sela, W. Song, H. Weatherspoon, R. van Renesse. In ACM TOCS, October 2017

Agenda

- The cloud
- Who created the cloud
- What is the cloud made
- Virtual machines
  - How far does the rabbit hole go?
  - Case studies

Takeaway

- The cloud provides an opportunity
  - Scale, performance, AI/ML
- Treat the cloud as a commodity
  - Sell virtual machines and give users illusion of owning their own machine
- Virtualization
  - Tradeoff performance and illusion of unlimited capability