CS 695: Virtualization and Cloud Computing

Lecture 1: Introduction

Mythili Vutukuru

IIT Bombay

Spring 2021

Virtualization and cloud computing

- What is the cloud?
  - Commodity servers with lots of compute and storage, connected with high speed networking, located in data centers
- What is virtualization?
  - Multiple virtual machines (VMs) can run inside a physical machine (PM)
  - VM gives user an illusion of running on a physical machine
  - Containers are like lightweight VMs
- Virtualization is a building block for cloud computing
  - Virtualization enables multiple clients share the cloud's compute resources
  - Multiple users on VMs/containers can share same cloud server
- In addition to compute, clouds also manage large amounts of data
  - Cloud storage/big data systems for efficient storage and retrieval of data

What is this course about?

- Two parts of the course:
  - Understanding how virtualization works
  - Basics of cloud storage frameworks

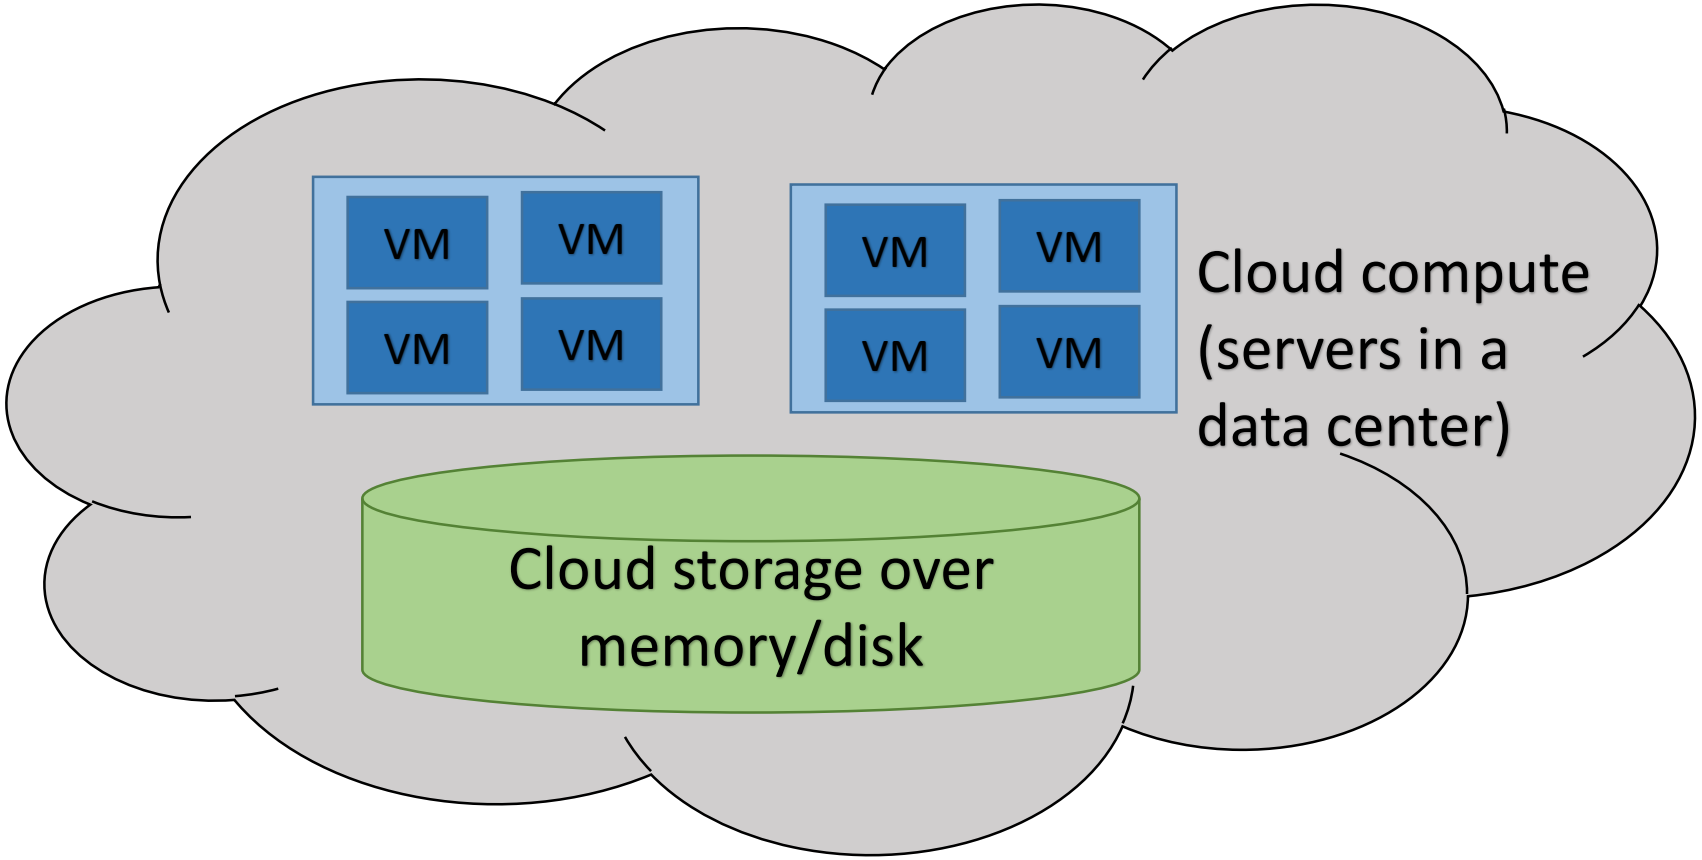

Why cloud computing?

- Public cloud providers (Amazon AWS, Microsoft Azure, Google Cloud etc) setup and maintain data centers with high-end servers
  - Powerful CPUs, lots of memory, disk storage etc., available to users
  - Organizations can also run a private cloud only for their users
- Why run applications on cloud and not on "bare metal" servers?
  - Multiplexing gains: multiple VMs can share the system resources
  - Lower overhead of maintenance: hardware/software maintained by providers
  - Flexibility: VMs can move to another machine if one fails
  - Pay as per usage: no need to invest in servers if only lightly used
- Disadvantages of running applications on cloud
  - Performance: longer delay to access servers via internet
  - Higher cost if heavily used

Virtualization terminology (1)

- We will study system virtualization, or how to run one full system (OS and applications) over another OS
  - We do not cover process virtualization (e.g., Java virtual machine) which lets a single process run on a different architecture from underlying machine
- Hypervisor or virtual machine monitor (VMM): a piece of software that allows multiple VMs to run on a physical machine (PM)
  - We will study how VMMs are designed

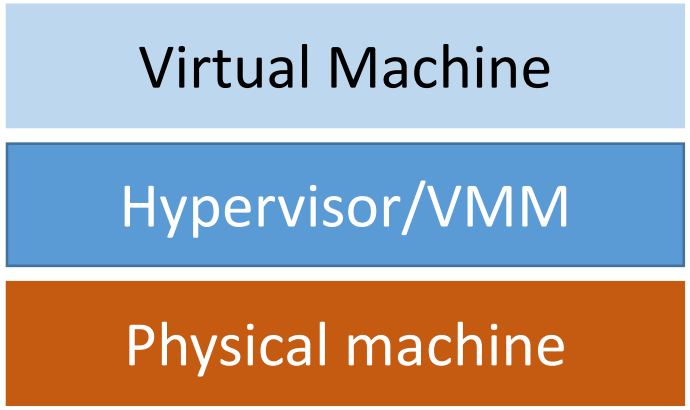

Virtualization terminology (2)

- Guest OS runs inside the VM, and host OS runs on the PM
- Type 1 hypervisor: runs directly on hardware, no need for host OS

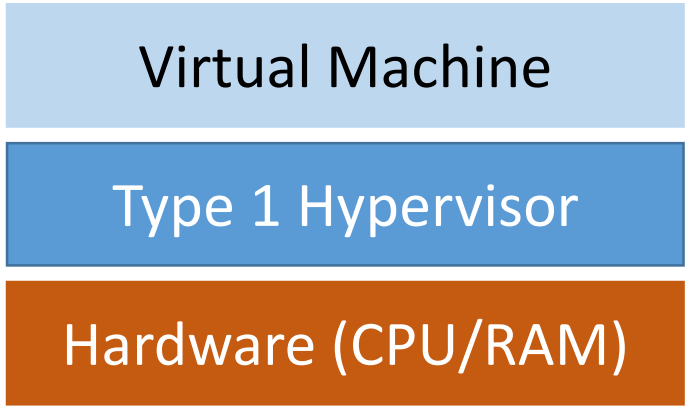

- Type 2 (hosted) hypervisor: runs as an application on top of host OS

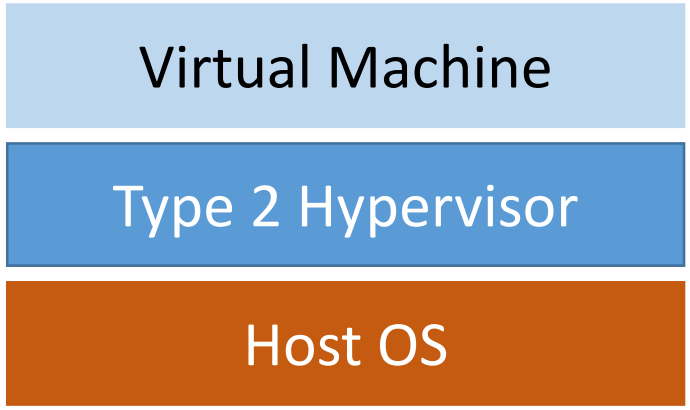

Challenges to virtualization

- Guest OS expects complete control over hardware, but VMM must multiplex multiple guests on the same hardware
  - Understand how operating systems work (prerequisite)
  - How to trick the guest OS into relinquishing hardware control?
- We will study the following ways to design virtual machine monitors
  - Hardware assisted virtualization (e.g., KVM/QEMU): modern CPUs have support for virtualization and VMMs are built over this support
  - Full virtualization (e.g., VMWare): Original technique to run unmodified OS over original hardware with no virtualization support
  - Paravirtualization (e.g., Xen): Modify OS source code to be compatible with virtualization
- Understand how CPU, memory, I/O devices are virtualized with each of the above techniques

Other topics in virtualization

- VM live migration and related ideas
  - VMs can moved from one physical machine to another
  - Why? Maintenance of machines in the cloud, fault tolerance etc.
  - How are VMs migrated without impacting the application in it?
  - Use similar techniques for other uses like VM checkpointing
- Containers: lightweight virtualization technique
  - Underlying Linux concepts of namespaces, Cgroups
  - Container frameworks like LXC, Docker, Kubernetes

Introduction to Cloud Computing

- Architecture of cloud applications: compute and storage options
  - Compute in VMs, containers etc.
  - Traditional storage in databases, now moving to simpler key-value stores etc.
- Cloud storage techniques
  - In-memory key-value stores (Amazon Dynamo)
  - Semi-structured data storage (Google Bigtable)
  - Application specific storage (Facebook's Haystack to store photos)
  - Caching-based optimizations to the cloud storage layer (Facebook's memcache)
- We will only briefly touch upon distributed systems concepts like replication, fault tolerance etc. as required
  - This is not a distributed systems course

Summary

- Introduction to virtualization and cloud computing
- Basic terminology and concepts
- Course outline
  - Techniques to design VMMs (hardware-assisted, full, para virtualization)
  - CPU, memory, I/O virtualization
  - VM live migration and related ideas
  - Containers
  - Cloud applications and cloud storage