

8.11.25

CS 695

lecture #1

Q. what is cloud computing?

- computing-as-a-service (x) (x)
- decoupled usage/access from implementation of functionality
- up/down scaling of service/resources...
(flexible/elastic resources) on-demand.
- pay-as-you-use / fine-grained metering
- (x) - reliability/security)...

service models

- IaaS - infrastructure compute infra.
cpu, mem, storage, ...
- PaaS - platform dev. platform, app engine
postman
- SaaS - software gmail, zotero, canva...
- FaaS - function/microservices
- *aaS - MLaaS
DBaaS
KVaaS

⑧ each of these service models need —

~ interface / APIs

~ abstractions — VMs, containers, etc. (design & impⁿ.)

~ mechanisms, policies

└ resource mgmt., provisioning, ...

⑨

user-view

- ~ on-demand & elastic
- ~ rapid / quick deployments
- ~ isolation, security
- ~ pay-on-use
- ~ performant

provider-view

- ~ demand ~ plan ~ increase
- ~ utilization / efficiency
- ~ features
- ~ minimize SLA violations
- ~ reduce costs

↑
(focus of the course)
for the most part

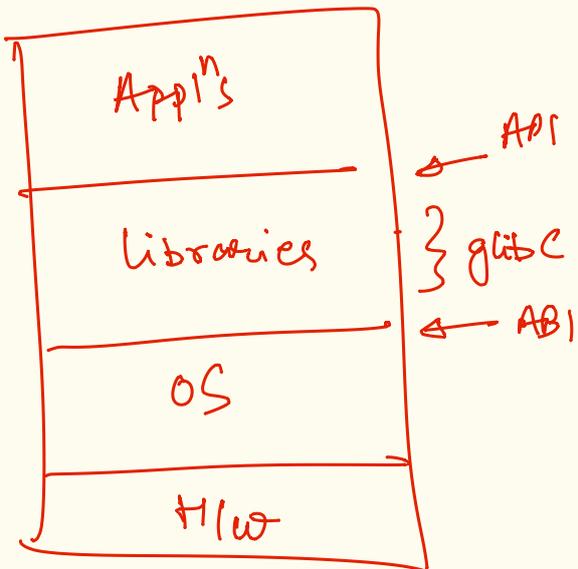
⑩

Q: what is this connected to virtualization?
why

~ secret sauce
of computing

ABSTRACTIONS

- + encapsulate functionality
- + provide interfaces
- decouple access / implementation



OS abstractions
 ~ process, files
 address space, ports, n/w
 locks/sync. primitives,
 n/w interfaces.

process-view
 on a m/c

- ~ address space of (zero starting - max)
- ~ cpu registers machine
- ~ files / directory structure
- ~ n/w endpoints
- ~ IP

OS-view of a machine

- ~ physical memory / addresses
- ~ interconnects (pci 3.0, pci 4.0, SATA, ...)
- ~ CPU type (32-64, x86, risc)
- ISA

↑ OS expects to provide abstractions.

↳ (machine) ←

⊛ virtualization / virtual machine

↳ build an abstraction for the OS-view.
 (virtual machine)