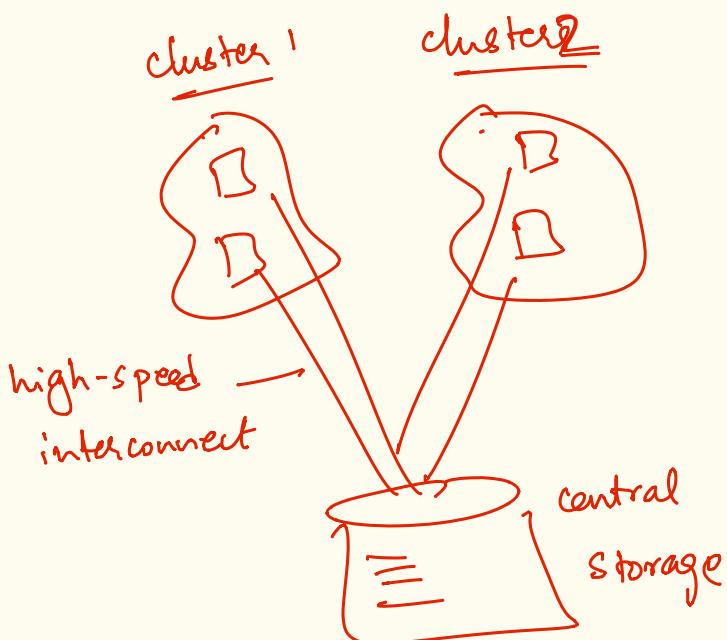


# CS 695

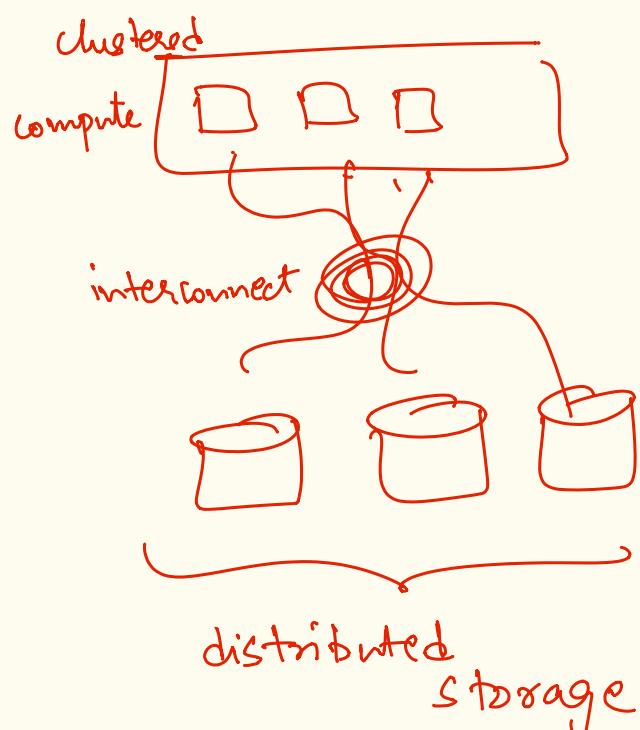
## ① data-centre architectures.

(a) clustered compute  
vs centralized storage



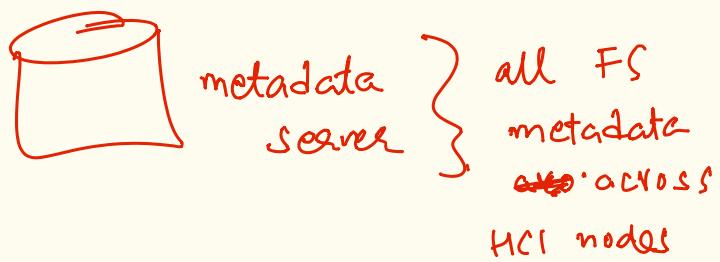
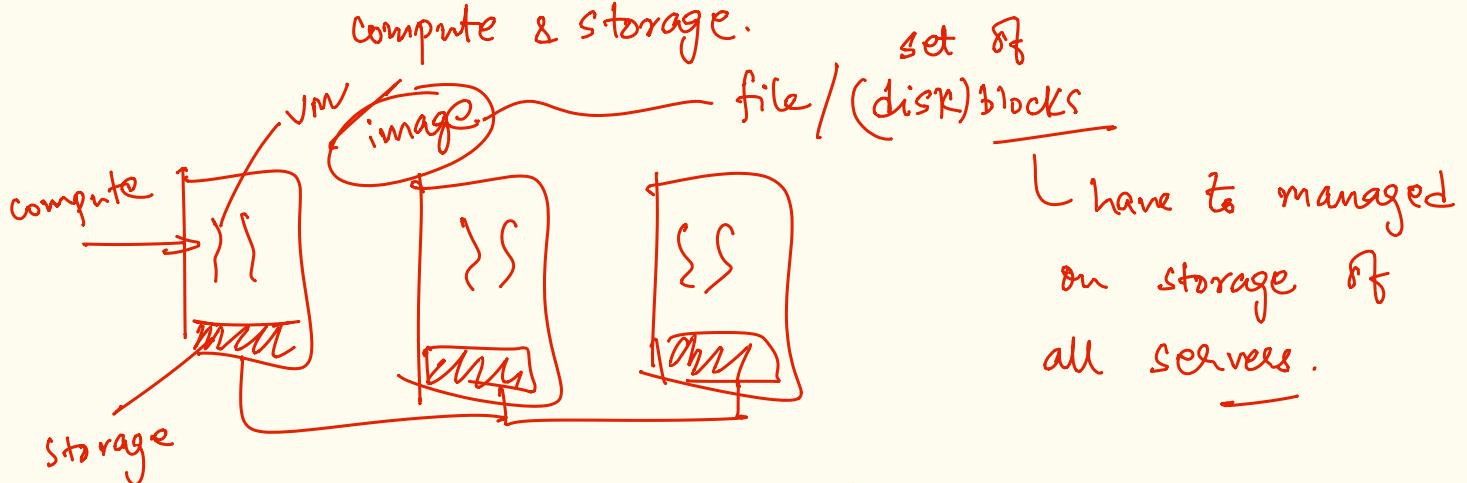
compute  
storage

(b) clustered compute  
clustered storage  
networked



(c) HCI - hyper converged infrastructure

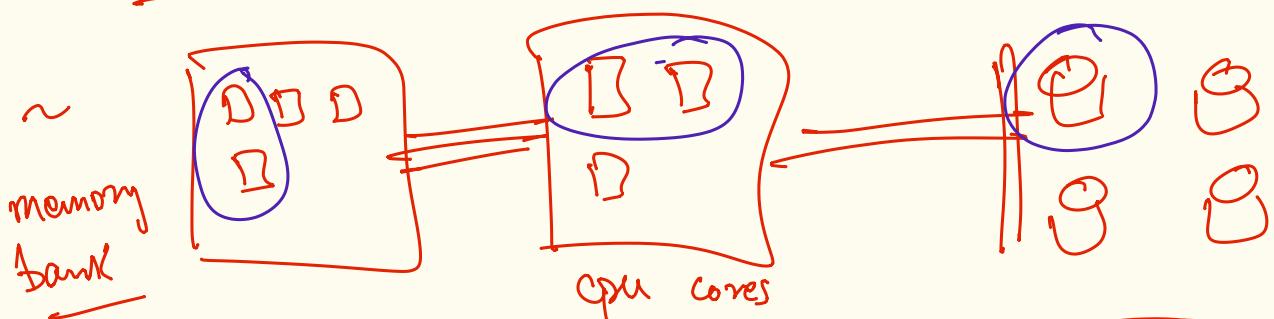
⇒ all nodes contribute to  
compute & storage.



## ② rack-scale architecture

~ disaggregate & ~~provision~~ provision & manage  
all / individual resources separately.

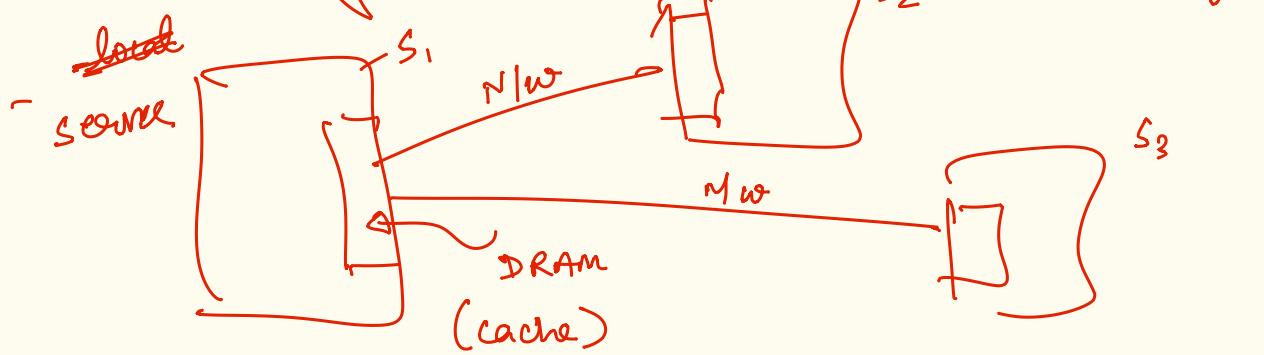
~ decouple all resource mgmt. & provisioning.



~ distributed OSes.

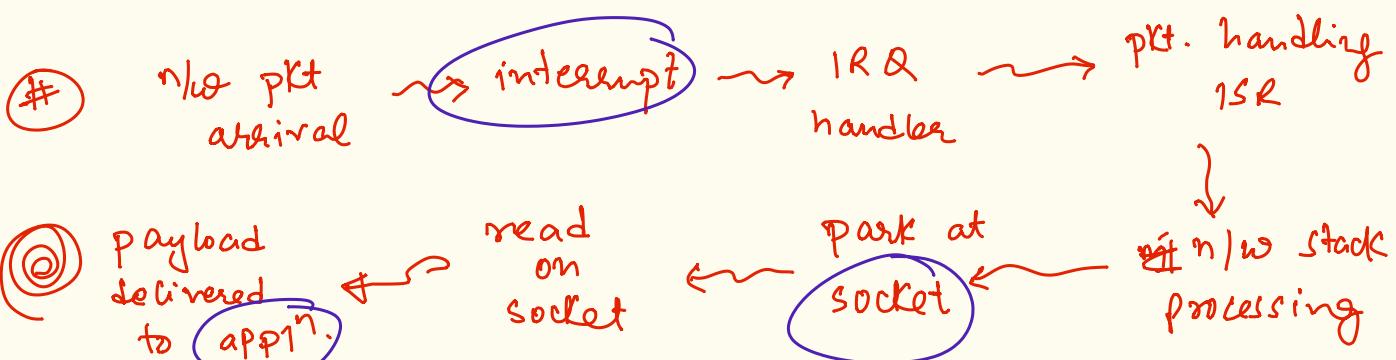
→ manage multiple machines like via single OS.

~ remote memory (RAM BUS)

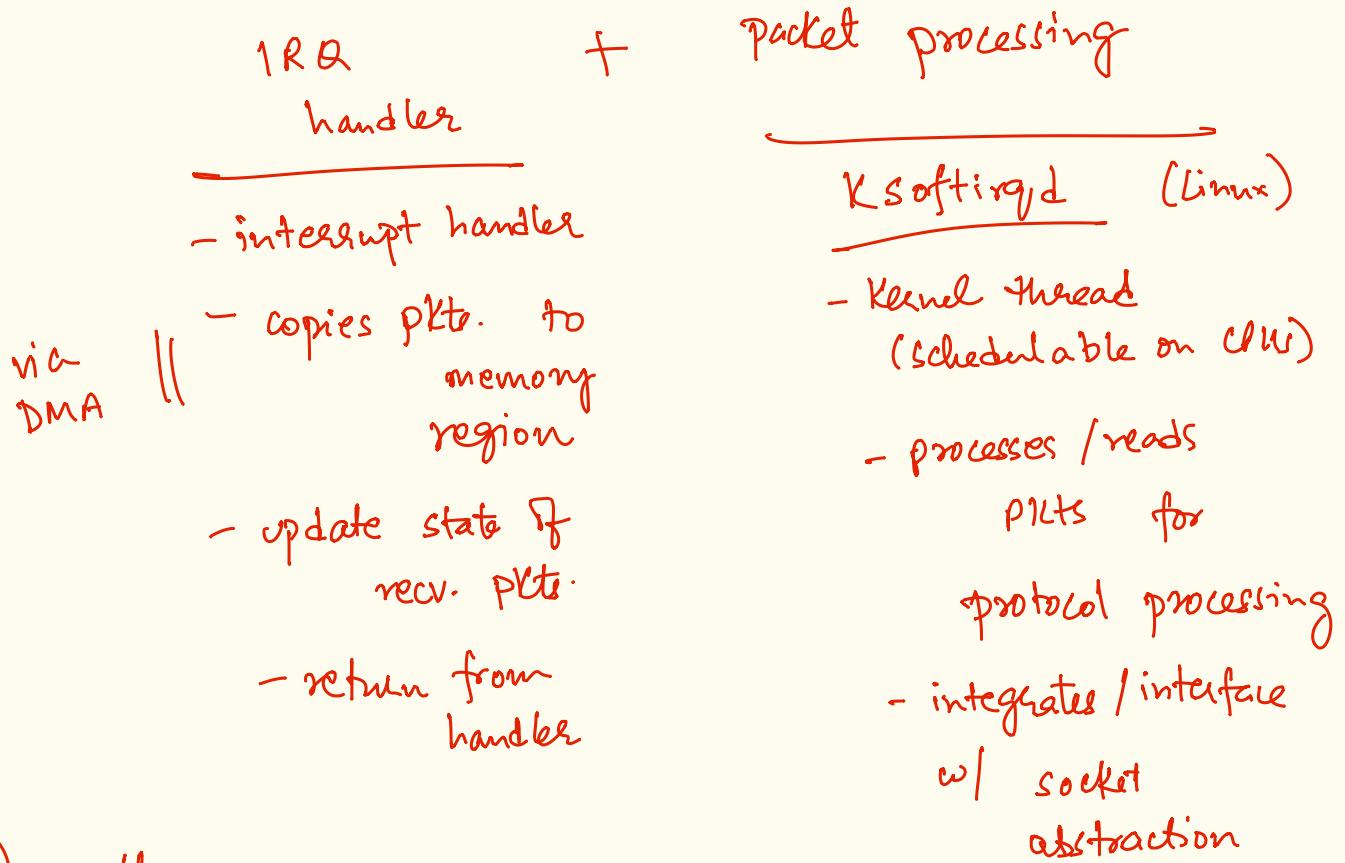


## ② networking in cloud-scale systems

~ n/w traffic / n/w processing is non-trivial.



### a) split-driver IO model.



### b) polling

blocking calls for read/write

poll / select / epoll

### c) get kernel out of critical path

#### Kernel bypass

DPDK ~ data plane development Kit

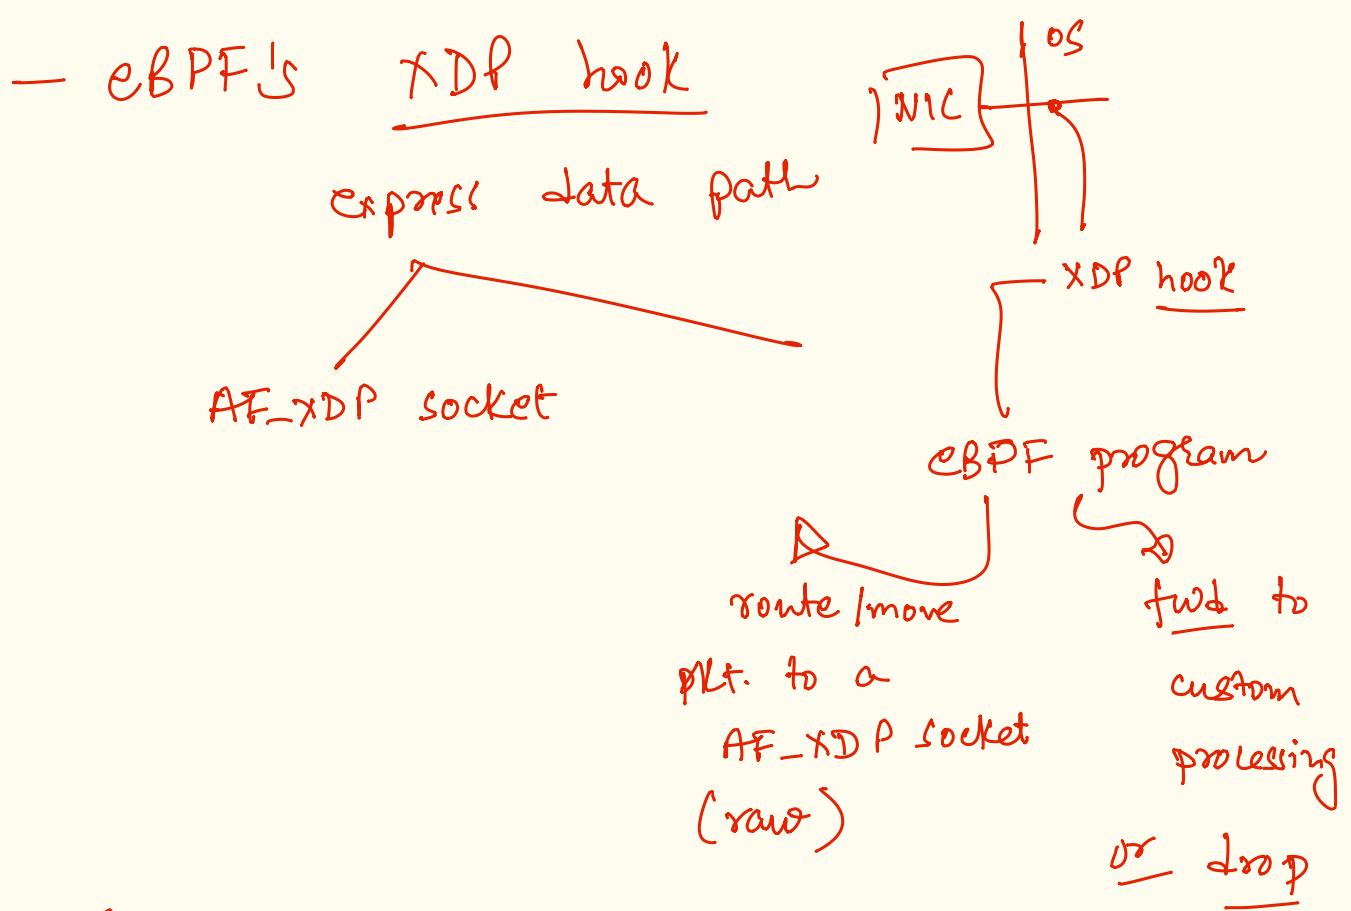
eg: firewalls

IDS

pkt. fingerprinting

NAT

} need access to either raw pkts. or custom header processing or both.



### ③ Storage

~ VMFS / GFS  
 ~ at-scale FS / storage solutions.

VMFS

{ image

composition of blocks .

{ blocks are duplicated

network storage

caching/ prefetching

deduplication

(\*)

## multi-cloud deployments

+ services across providers.

(\*)

## acceleration-as-service

+ via GPUs.

at scale mgmt. of GPU provisioning  
& usage.

cloud



we are probably here!  
(via CS695)