Choosing Hardware

Cloud and Virtualization Workshop

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

CPU

- All modern CPUs are 64-bit ("x86_64", "amd64")
- All modern CPUs should have hardware virtualization support
  - Intel "VT-x" & "VT-d"; AMD "AMD-V" & "AMD-Vi"
  - You may need to enable it in the BIOS
- Sufficient cores
  - although sharing cores between VMs is fine if workloads are intermittent
  - core types, like Intel's performance (p) and efficient (e) cores
- A dual-socket motherboard lets you add another CPU later

RAM

- Each VM guest is assigned a certain amount of RAM; make sure you have enough RAM in total
  - Host swapping to disk is a no-no
- Some clever tweaks possible
  - e.g. Linux ksmd; kernel shared memory daemon
  - Only useful if running many *identical* VMs
- Far better not to overcommit your RAM in the first place
- Allow some extra for host OS and overheads (e.g. linstor, ceph)
- Prefer systems which support ECC (parity) RAM (for servers)
  - DDR5 DIMMs have "on-die ECC" but hide its operation from the host

Disk

- A hard drive is "spinning rust"
- SSDs are much faster
  - Look for Enterprise drives (longer life under load)
- RAID5/RAID6 performs very poorly for VMs
  - very poor write speeds
  - very poor read and write speeds when degraded
  - only use for "cold" archival data
- Whole separate presentation on this

Network bandwidth

- 1Gbps = 100MB/sec = about the bandwidth of one hard drive
- One SSD typically 500MB/sec or more
- Use a separate NIC on servers for storage!
- Ideally 10Gbps or 25Gbps
- Consider tuning MTU=9000 ("jumbo frames") for storage network
  - however this can cause outages unless you're sure you've configured all your switches correctly

Other recommended features

- Multiple NICs
  - 10G+: VM frontend, Storage
  - 1G+: management, cluster sync
  - Consider bonded pairs of NICs for redundancy / load sharing
- Dual power supplies
  - Connected to separate power bars fed from separate UPSes
- "Lights out management" for remote server control
  - e.g. power on/off, reinstall OS, monitor chassis temperature…
  - Many manufacturers charge for full remote VGA console (IP KVM)
  - But may offer a cheaper IPMI option with remote serial port access

OOB / LOM / IPMI / iDRAC

- **OOB**: Out of Band (separate management network)
- **LOM**: Lights Out Management
  - **ILO**: Integrated Lights Out Management
- **IPMI**: Intelligent Platform Management Interface (standard)
  - Standard clients like ipmitool, ipmiutil
- **iDRAC**: Integrated Dell Remote Access Controller
- … Many vendors have their commercial equivalents

Sample Hardware

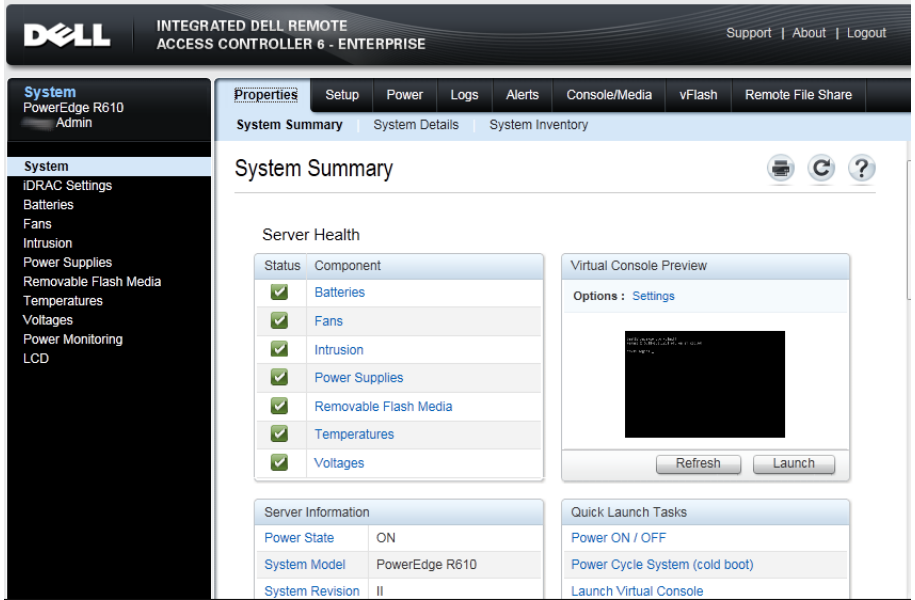

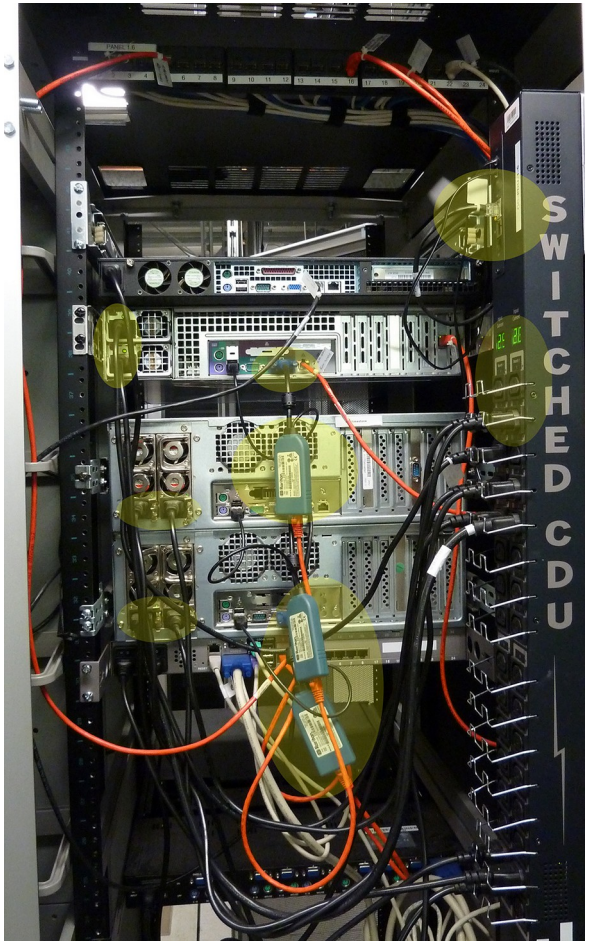

Summary – choosing hardware

- Choose servers with VT-(x,d) or AMD-(V,Vi) support
- Buy enough RAM for all your VMs combined
  - and spare DIMM slots for expansion
- Install multiple hard drives
  - but don't use RAID5 or RAID6 for VM images
- Install multiple NICs
  - or expansion slots to add them later
- Consider some out-of-band management option