Farm/Cluster Observations

Thomas A. Yanuklis

May 23, 2001

Machine Life-cycles Notes

- At the RCF machines have a life span of 3 years before obsolescence.
- The specification for machines also becomes more specific as experimenters' code base and usage model become more understood.
- The kind of machine has changed a bit, the systems are now ordered with proportionally more memory and local storage than before, due to a slight shift in the processing model(and a little code bloat).
- On the software end, we are pretty conservative, the farm runs RedHat 6.1 officially, and we are currently testing the feasibility of running RedHat 7.1.
- The evolution of processors, with the help of Moore's law keeps making buying more SPECint95s cheaper over time.

Operational Notes

- We do not benchmark the machines; we give machines that are on loan from vendors to the users, and ask the users for their comments on them.
- There are only two type of machines on the RCF farm, interactive and batch. The delineation between the two is minor; machines can be converted between either mode of operation.
- Our purchasing criteria is price performance, we want inexpensive SPECint95s, that meet our criteria for management, and that are compatible with our installation image.

Unforeseen Consequences

- Machines produce heat and need to be cooled, the need to purchase extra cooling capacity becomes a necessity as a farm expands.
- Power is also another hidden gotcha, as rack density has increased the need to Higher Voltage power supplies develops.
- Space, where to put all the systems actually seems to be less of a problem given greater rack density, the tradeoff is a greater need of cooling and power.
- The users of the system may shift their processing model, is your organization able to re-tool and meet the demand?
- Will Moore’s law cease to be accurate in the near future?
- How will vendor changes to a product line effect a farm?
- How often will you upgrade the software on the farm, and who does it, who approves the changes?

Physical Characteristics of RCF Systems

- Dual/Multi Processor System (x86 architecture)
- On-board Network support[a]
- Minimum 512Mb memory[b]
- Minimum 36Gb Disk[c]
- Floppy Disk Drive[d]
- CD-ROM[e]

---

[a] Well, at least a Fast Ethernet NIC of some type.

[b] This number is subject to change, it's a moving target.

[c] ibid.

[d] Floppy disk drives have proved to be a necessity when diagnosing some problems.

[e] We have purchased systems w/o CD-ROM, but have latter regretted it.

Management Requirements of Systems

- The ability to install operating system images on the machines (e.g. RedHat's kickstart installer, or VA's Systemimager).
- The abilities to remotely power-on,power-off, and reboot nodes are required.
- The ability to remotely monitor the console is preferred.
- Out Of Band Management from Vendors:
  - VACM Controls nodes via RocketPort Serial multiplexer in older machines.
  - VA's new system uses a Cable chain to a central controller unit that communicates w/ the mayor node.
  - IBM uses a Cable Chain System, that uses "System Management Processors", PCI cards with onboard Ethernet, commands are issues to the cards over a private subnetwork.