USC HPCC

Sun HPTC Meeting

Current 13.81TF Cluster

Today and Tomorrow

HPCC

- Provide common facilities and services for a large cross section of the university that requires leading edge computational and networking resources.
- Leverage USC central resources with externally funded projects.
- Campus ISD support for base facilities
- Not just usual suspects play
- New organization is evolving at core of USC

HPCC

- Faculty advisory group
- Allocation committee
  - Large Disk/cpu
  - Favors larger scale problems.
  - Favors interdisciplinary proposals
  - 5 users over 500k hrs
  - Routinely setup 512 node runs
  - 256 are standard queue
- Large memory jobs are also common
- Viz/Data initiatives
- 13.8TF cluster

HPCC

- Highlights
  - Leverage across university.
  - ISD is catalyst for inter-disciplinary work
    - USC strategic plan stresses this
  - 3,000,000+ node hours in last year
  - 8 ‘condo’ users 600+ nodes
    - Newest was 256 of 512 4100s
  - 70TB ‘condo’ disk
  - Networking
    - Los Nettos
    - Pacwave
    - Dark fiber/waves
  - Shared staff
  - Fits in NSF layer model for cyber-infrastructure

Current HPCC Resources

- High Performance Computing Resources
  - Linux Cluster (1830 nodes/5384cpus, 2Gb/sec Myrinet), many file head nodes
    - ~100TB shared disk, 18GB - 40GB local disk per node.
    - Ranked in #2 for USA academic clusters last time. (stay tuned)
  - Myrinet switch is 1952 total nodes possible now.
  - Adding nodes funded by USC research groups.
  - Sun Core Servers (E15k shared memory)
    - 72 processors, 288GB memory, 30TB shared disk
    - Looking to large core count opteron possible replacement
  - Mass Storage Facilities (QFS)
    - 1.1PB on tape

Building a Big Cluster

- Large cluster represent unique challenges
  - Power
  - A/c
    - Air flow
      - Hot spots
      - Volume
    - What happens when a/c fails
  - Wiring
    - Density
    - Testing
    - Power cabling.
    - Blocking cooling
  - Going to new model in our new facility
    - Increase in floor depth
      - 2 ft floor height
    - 60% overhead air flow
      - 8 foot square duct in ceiling
      - Very high ceiling allows this

Why It's Hard

- Large cluster represent unique challenges
  - Software installation
    - Non-homogenous cluster
      - Built over time
      - Different vendors
      - Different hardware base configurations
  - Merging/adding new 'chunks' is complicated
    - High speed network (Myrinet)
      - High density spine
      - Not full mesh (cost tradeoffs)
    - Gb ethernet
      - 10Gb interconnect between switches
      - 1Gb to each host
    - How to do this in running cluster
      - VERY carefully
      - Pre-position cables/switches
      - Lots of labor in short time

Configuration

- Types of processors
  - Sun V60 (3.0Ghz, 500Hhz, 2GB) (some 4GB)
  - Dell 3.2 (3.2Ghz, 800Mhz, 2GB)
  - Sun V20z (opterons) (2.2Ghz, 2GB)
  - Sun V20z (opterons) (2.0Ghz, 4GB dual core, dual node)
  - Sun 4100 (opterons) (2.0Ghz, 4GB dual core, dual node)
  - Able to mix arch.
- Goal is to increase in cores/processors per node
  - 8 cores in next 12-24 months
  - All nodes 64 bit in next year
- Interesting power trade-offs with opterons vrs xeon

Configuration

- Networks
  - Myrinet
  - 7 spine pairs on one chassis 128 ports each 896 total potential
    - New 4/1 cards 4 fiber on each port
    - Reduces clutter
    - Room for 3 more spine pairs
  - 15 128 port edge switches
    - 96 host ports
    - 32 spine ports
    - 3/1 host/spine ratio
  - 2 256 port edge switches
    - 256 host ports
    - 128 spine ports
    - 2/1 host/edge ratio
  - Ethernet for each node
    - File I/O
    - Many back-end file servers (15k, 440s,240s,v20zs,v40zs (qfs fses, nfs access)
  - Console concentrators for all nodes (management)

Hot Stuff

- Power loads at full song
  - 480v transformers
    - Another 80KW on 208v (runs building lights and overhead a/c (lieberts) and one main a/c
  - We have surpassed power available on our generator
    - Special setup to allow cluster to be off generator
  - Slightly less than with 2.8Ghz xeons.

Transformer	Before	During
TI	290kw	350kw
T2	400kw	750kw

Benchmark

- This one was very hard
- 4100s changed paradigm
  - Fan speed
  - More heat leakage
- Noticed issues with previous v20zs
  - Not seen in normal operations
  - 4-5 nodes had dry(ing) thermal paste
  - Showed up most on nodes near top of racks
  - They got more heat leakage due to increase air speed
- Gave back maybe 5-10% not sure
- Had to be very aggressive with heat mitigation
  - Lots of air dams
- Next generation run will be in new facility (discuss later)
  - Drawing and special mitigations in cluster room in new faciltiy

Benchmark Details

- Started with 1/2 memory on 5376 procs (512MB/proc)
  - N=500000, 72x74 (5328, took out top nodes in several racks)
  - Ran this 4 times without completion
  - Mostly v20zs with heat issues (>70C shutdowns)
  - Setup monitor process for cpu heat status
  - Stole cardboard from homeless people
  - Moves some tiles near end of rows.
  - Complete run at N=500000
    - 13.81TF
    - 65% efficiency
    - 6033 seconds
    - Most cpus were in 65C range one near shutdown
- Max memory would take 308 mins
  - 2/3n**3+o(n**2) increase in run time for same ‘speed’
  - ‘No way’ also ran out of downtime

Real Life Cluster in Datacenter

- Slide show.
  - Heat density
  - Power
    - Lets melt
  - Building/logistics

New Data Center

- Online in fall
- Some blueprints