

# Setting Up and Using a Small Linux Cluster

J. David Moulton<sup>†</sup>

(moulton@lanl.gov)

Ben Bergen<sup>†</sup> , Markus Berndt<sup>†</sup> , Patrick Goda<sup>‡,†</sup> , Michael Warren<sup>‡</sup>  
(bergen@t7.lanl.gov, berndt@lanl.gov, pgoda@lanl.gov, msw@lanl.gov)

<sup>†</sup>Mathematical Modeling and Analysis Group

<sup>‡</sup>Theoretical Astrophysics

Los Alamos National Laboratory  
Los Alamos, NM

July 8, 2001.

This research was performed under the auspices of the U.S. Department of Energy, under contract W-7405-ENG-36, and is to be referenced as LA-UR 01-3854. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.



# Beowulf in a Day



- *Introduction* 9:00 – 10:00
  - hardware overview ( memory, cpu, network etc. )
  - purchasing guidelines and *hands on* node assembly
  - ▣▣▣▣▶ *Coffee Break* 10:00 – 10:30
- *Installing and Configuring Linux* 10:30 – 12:30
  - RedHat GUI installation and automation with kickstart
  - configuration and kickstart customization
  - ▣▣▣▣▶ *Lunch Break* 12:30 – 2:00
- *Development Tools and Software* 2:00 – 3:00
  - installing MPI (mpich, lam)
  - benchmarking, development and debugging
  - ▣▣▣▣▶ *Coffee Break* 3:00 – 3:30
- *Real World Codes on Linux Clusters* 3:30 – 4:30
  - highlights from 3 years of research on Avalon
  - cutting edge hardware (Gigabit, RAID etc. )
- *Concluding Remarks and Open Discussion* 4:30 – 5:00



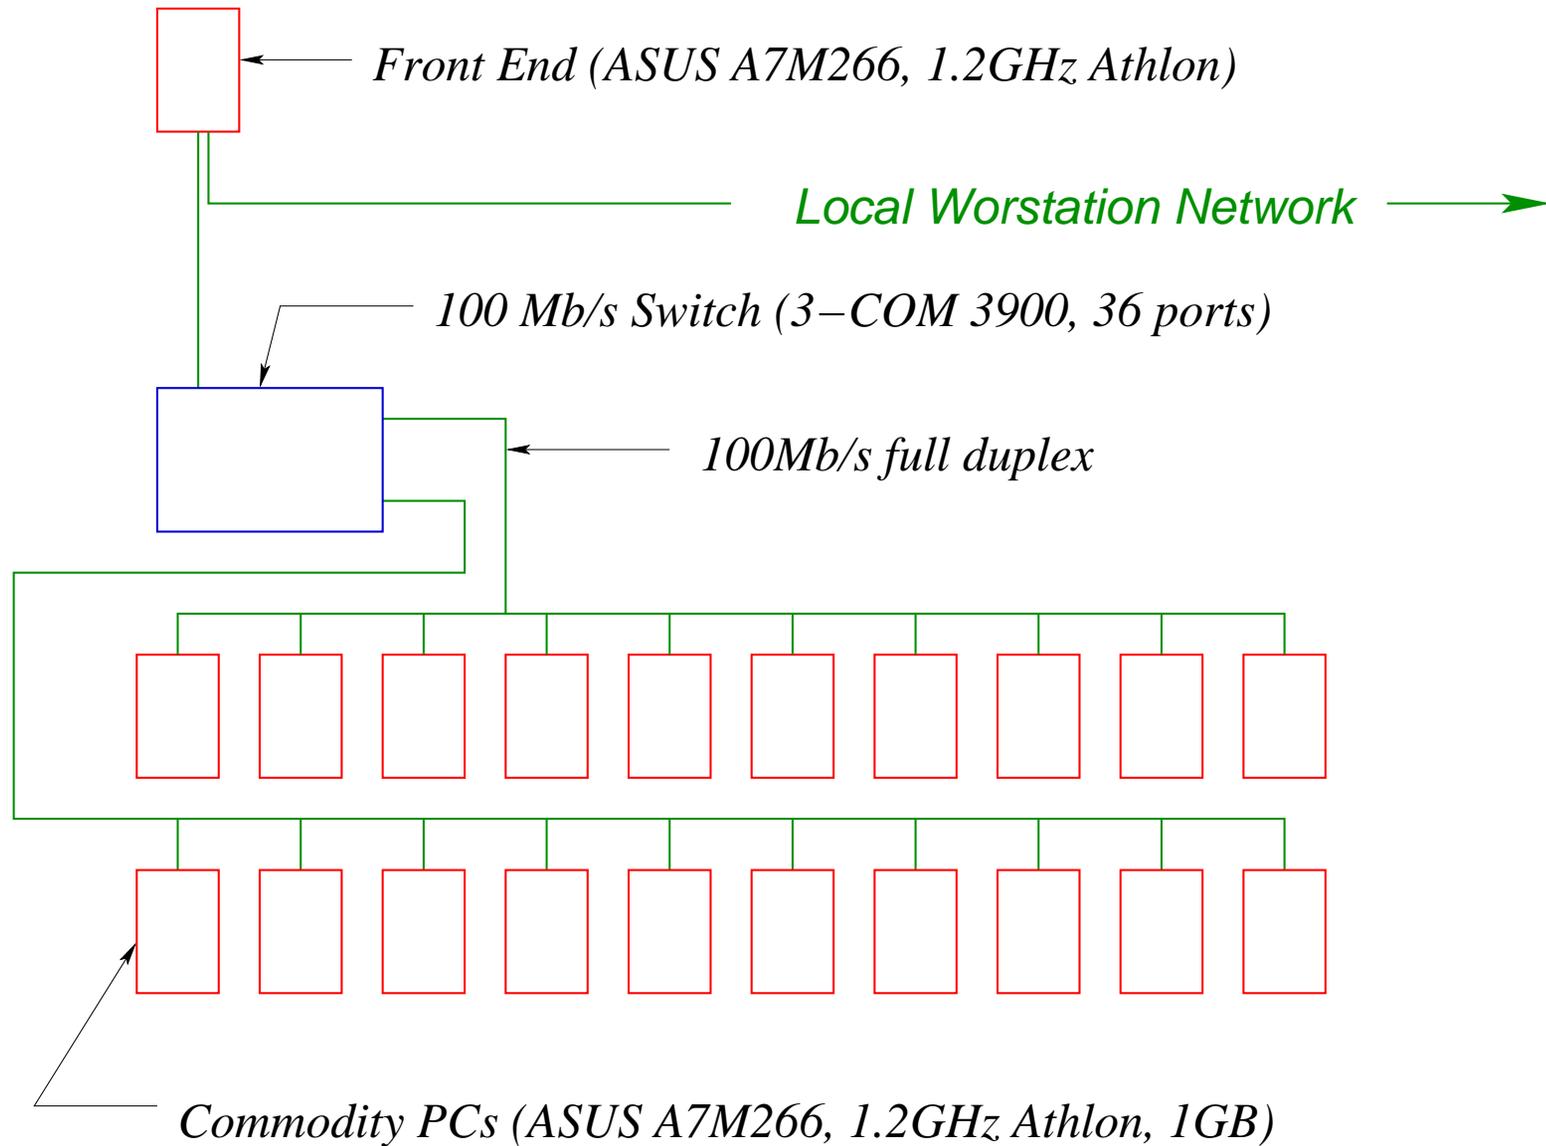
# Outline: Cluster Hardware Selection



- *Beowulf Clusters*
  - a brief history
  - schematics and pictures
- *Commodity System Components*
  - system boards
  - CPUs: Athlon, Intel PIII and P4 etc.
  - memory: PC133 SDRAM, PC2100 DDR etc.
  - network cards and switches
  - miscellaneous: disks, video cards, cases
- *Purchasing*
  - online shopping
  - a full system vs. parts for assembly
- *Hands on node (dis)assembly*



# A Beowulf-Class Cluster Schematic





# Desktop Cases vs. Rackmount



*A shelf of Avalon with 40 desktop systems*



*1U rackmount systems*



# Commodity CPUs



## *Pentium III (Coppermine) 1.0GHz*

- 133 MHz front side bus  $\Rightarrow$  memory bandwidth  $\approx$  1.04GB/s
- 256KB of full speed L2 cache
- single floating point pipeline

## *Athlon 1.2GHz*

- 266MHz front side bus  $\Rightarrow$  memory bandwidth  $\approx$  2.08GB/s
- 128KB split L1 cache (instruction and data)
- 2 floating point pipelines (separate add and multiply through 3D now )

## *Pentium 4 (Xeon) 1.7GHz*

- 400 MHz front side bus  $\Rightarrow$  memory bandwidth  $\approx$  3.1GB/s
- 256KB of full speed L2 cache
- single floating point pipeline (more stages)



# Memory



## *SDRAM (Pentium III)*

- Synchronous Dynamic Random Access Memory (SDRAM)
- PC133 SDRAM is the specification for 133MHz FSB
- only a few parameters, CAS-2 is best

## *DDR SDRAM (Athlon)*

- Double Data Rate (DDR) SDRAM  
(see [Tom's Hardware: DDR-SDRAM Has Finally Arrived](#))
- PC2100 SDRAM is the specification for 266MHz FSB
- latency is similar to SDRAM

## *RAMBUS (Pentium 4)*

- new memory design based on a bus architecture
- PC800 RIMM (RAMBUS Inline Memory Module) for 400MHz FSB
- higher bandwidth but worse latency



# Networking



## *Fast Ethernet (100BaseT)*

- 100 Mb/s NICs are incredibly affordable (\$25 ea.)
- 100 BaseT switches are very affordable (\$10 per port)
- drivers are well developed and bonding works well.

## *Gigabit Ethernet (1000BaseT)*

- Gigabit NICs are dropping in price (\$130 ea.)
- On the Athlon's 32bit/33MHz PCI Bandwidth is  $\approx 675\text{Mb/s}$
- Latency is higher than Fast Ethernet
- switches are still expensive and have few ports (e.g., 16).

## *Other Gigabit solutions*

- Myrinet is very costly, in general, not really commodity hardware



## Miscellaneous



### *Disk Drive:*

- ATA-100 is inexpensive, IBM Deskstar disks are excellent
- SCSI isn't likely to be worth the extra cost

### *Extras:*

- AGP video on the front end and/or workstations
- preferably builtin video on the nodes
- cases: Antec etc.



## Purchasing: Do Your Homework



*Use online resources for technical information and pricing*

- use search engines, particularly Google! ( [www.google.com](http://www.google.com) )
- hardware review sites are a big help, but tend to be geared at gamers (e.g., [www.tomshardware.com](http://www.tomshardware.com) )
- hardware vendor sites usually have technical information
- virtual comparison shopping (e.g, [www.shopper.com](http://www.shopper.com) and [www.pricewatch.com](http://www.pricewatch.com) )

*Don't get too close to the cutting edge*

The premium you pay may only buy you frustration

*The 10-15% more you pay to buy preassembled may be worth it!*

*Prototype, the machine(s) will be useful somewhere.*