

# The Case for NOW

January 29, 2004

## I. Background

Huge effort in parallel computing throughout early 90's. It was the "P2P" of its time, with new conferences, lots of top researchers, and lots of funding...

Key goals:

- o scalability, raw FLOPS (how many CPUs)
- o network bandwidth
- o programming languages (parallel fortran)
- o shared memory models

But some big problems:

- o workstations had much better single CPU performance -- mostly due to 18 month lag in usage of a particular processor
- o very high engineering cost spread over relatively few machines (sold) => high cost per CPU for these machines to amortize the R&D
- o awkward development environments
- o custom OS was also behind in features/reliability
- o hard to upgrade
- o somewhat less reliable than workstations
- o hard to program even with some language help
- o various PhDs on all of these topics...

Only real advantage of the large parallel machine: backplane bandwidth

- o solution: create a cluster network that has similar bandwidth (if not latency)
- o this was eventually fixed by AM work at Berkeley and the U-Net work at Cornell (and also partially in the VI interface for Windows).

Clusters:

- o key idea: have to reuse the whole box, not just the CPU
- o implies: better performance, much better cost, latest OS and tools
- o challenges: even harder to program, network still not as good
- o extra challenge: can you make use of idle workstations? (more on this later)
- o huge amount of aggregate disk I/O (and seeks)
- o software RAID (rather than hardware) -- this took many years to really work

### Big Issues:

- o how do you program a cluster?
- o how do you deal with partial failure? (you how potentially have all of the problems of distributed systems!)
- o how to get a global system view: scheduling, file systems (easy), shared caching, namespaces?

### Killer App?

- o turned out to be web servers, led by Inktomi work in particular
- o Advantages for web servers: incremental scalability, fault tolerance, cost
- o Eventually worked well for traditional science applications as well, but mostly for those without the need for fine-grain cluster-wide sharing. E.g. rendering works great, but sparse matrix apps are much harder

### Things that didn't work out as planned:

- o no real use of idle workstations -- there are some counterexamples, mostly in graphics and simulation (and SETI !). Machine cost is minor compared to other costs and complexities.
- o Winner wasn't workstation vendors: really PC vendors plus Linux (see Beowulf project for example)
- o network RAM has never really made it big. Possibly due to security?
- o not much use of software fault isolation for the GLUnix layer. More a traditional layer under processes.
- o security? in practice it is provided by physical security for big clusters (not by software on users' desktops)