

Clustering for Research and Production Scale, Parallel and Distributed Computing

*Dr. Anthony Skjellum, President, MPI Software Technology Inc., USA
tony@mpi-softtech.com*

Abstract: A lot of attention has been paid to so-called Beowulf/Avalon clusters, where PCs or Alphas are strung together with 100Mbit/s Ethernet and portable programs from supercomputers have been run on these, particularly when modest bandwidth and latency requirements are posed by the example applications. In addition, heroic efforts to scale clusters using early gigabit/s scalable fabrics has been done across the world, but these systems, like the Beowulf counterparts, have relied on software from the previous generation of multicomputers and supercomputing systems.

However, commercial-grade software tools (middleware and distributed environments) for clusters have matured considerably since the initial Beowulf type experiments, as have the availability of easy-to-use cluster interconnects. In this talk we review the technical achievements thus far in production-grade environments for both message passing and cluster scheduling, both for NT and Linux. This talk emphasizes the option of having tools and hardware that is scalable, to varying degrees, and presents a taxonomy of hardware, software, and applications that divides the space of activities and also seeks to establish areas where additional opportunities for new software and other tools exist. Issues of security and scalability are considered as are cost of ownership vs. freeware, with proposed economic models for both company and university adopters of clusters. We discuss pros and cons of open source vs. commercial products as viable options moving forward.