

**Scalable Cache
Coherent Shared
Memory at
Cluster Prices**



White Paper

SMP Redux: You *Can* Have It All

Large scale plug-and-play SMP delivers mainframe performance at commodity prices

By: Douglas Eadline

About the Author: Douglas Eadline, Ph.D. is the Senior HPC Editor for Linux Magazine and has been writing and working with high performance computers and Linux for over twenty years. He is also the Editor of Cluster-Monkey.net.

NW3V1

ABSTRACT

When High Performance Computing (HPC) is mentioned, one often envisions a large expensive mainframe built by a company like Cray, IBM, SGI, or Sun (now Oracle). These machines are highly parallel and employ a Symmetric Multi-Processing (SMP) design that provides global memory and process spaces. Delivering this level of performance has always been expensive due to the custom engineering required to integrate a large number of processors into a shared memory environment. Indeed, the added expense of large-scale SMP systems has pushed the market to use a “cluster approach,” where a large number of commodity server machines connected together are used as a single resource. It is well known that clusters are harder to manage and less efficient than SMP mainframe systems.

The cost of large-scale SMP systems has been an impediment to their widespread adoption, when compared to the cost of commodity server clusters. In this paper we will contrast current SMP and cluster designs in the context of HPC. We will also introduce a breakthrough technology from Numascale that allows commodity hardware to be combined into a cost-effective and scalable SMP system that delivers the best of both worlds – the low cost of clusters and the efficiency of SMP.