Abstract

The development of energy-efficient server clusters requires the study of different request dispatch policies applied by the central access point of the cluster, and the application of hardware techniques that provide for the best usage of server resources. However, energy efficiency should not be attained at the expense of a poor quality of service. The work proposed here describes a technique that tries to balance energy consumption and adequate application response times in heterogeneous server clusters. This proposal was evaluated through simulation and showed results that outperform state-of-the-art techniques.