

Proposal to the Consortium for Embedded and Internetworking Technologies

Project Title: A Case for an Inexpensive, Highly Available iSCSI Storage Solution

(This is an extension of our 2002 proposal
“Development of an iSCSI Storage Manager with Virtualization
for eLinux on Intel 80321 Processor”)

Arun Sen and Andrea Richa
Computer Science and Engineering

Abstract

The objective of this research project is the design and implementation of an *inexpensive, highly available iSCSI Storage Solution*. The target embedded environment for this project is a low cost I/O processor such as the Intel 80321 or the Motorola MPC8240. The proposed project is an extension of our ongoing project on the development of an iSCSI Storage Manager with Virtualization for eLinux on the Intel 80321 Processor. The current project is scheduled to end in June 2003. This research will leverage the knowledge and the experience gained through the current project, to develop an inexpensive, highly available storage solution.

In our current project, we have implemented file-system level virtualization, which allows hosts to share and use remote data located anywhere on the Internet, through a distributed file system. The virtualization engine acts as a NAS server. It provides a storage-level abstraction for remote iSCSI targets (storage devices), by virtualizing them into volumes of configurable sizes. In addition to this virtualization task, the current research is also examining the issues of *Distributed Storage Caching* in a multiple storage manager environment. As a part of the current project, a significant amount of performance data for the storage architecture has been collected and we are currently in the process of analyzing such data.

The proposed research plans to extend our ongoing project by incorporating *reliability, availability and performance* to our current architecture. Using the knowledge and the data obtained through our current project, we plan to enhance the current architecture into a *large scale, distributed, highly available system*. In the new environment, Storage Management and Virtualization will be *distributed among multiple virtualization engines*. The engines will communicate among themselves to provide highly available access to the remote SCSI devices. The virtualization engines and the SCSI devices would communicate via the iSCSI protocol. The hosts would access the remote data through the Storage Manager using some distributed file system. Distributed Storage Caching and load balancing techniques will be utilized to improve the overall system performance. Failure recovery mechanisms that are transparent to the hosts will also be incorporated. The primary outcome of the project will be a *distributed, highly available storage solution* with low *Total Cost of Ownership (TCO)*. It is envisaged that such a solution will prove very effective for the middle to high range of enterprise storage applications.