Problem Statement

The goal of this project is to design a clustering system that addresses the basic issues of small organizations. They require a system that is inexpensive, reliable, easy to install, easy to maintain, and that has very little user (admin) interaction. They also require that basic services such ftp, file serving, and email will run on the cluster as seamlessly as they would on a single machine.

We aim to address the easy installation issue by writing turnkey installation code that automatically installs the required components of our system on each of two machines with little to no user prompting. We will setup one system to start up as an ALPHA server, the main system that offers services to the users. The second will be configured to start up as a BETA, which must take over the services if the ALPHA ever fails. The user will not be able to tell which of these systems is providing the services, as it will appear as one system to computers outside of the cluster.

We will provide an easy to use graphical user interface (GUI) for the person administering the cluster. This GUI will have a stoplight-type of alert which visually informs the administrator of the health of the cluster. Some key factors shown by this alert are: which server is currently ALPHA, whether or not the Heartbeat Protocol working properly, etc. The GUI will also have a services tree, which dynamically updates the status of the ftp, file serving, mail and other services running on the cluster. The GUI will also allow the administrator to

setup users authorized to use the ftp and mail services, the IP addresses used within the cluster, and to view the system logs.

We look to make the system reliable by taking widely tested open source clustering code and tailoring it for a cluster for a small organization. The benefit of using open source is that no royalties are required for use of the code, hence reducing the cost of development to next-to-nothing. We will be building our system using the open-source operating system, Linux, as this will ensure that whatever code we build upon will have already been widely tested in many environments and situations. Linux being an open source project is also well documented with a large number of people able to answer questions or fix problems related to Linux. With all of this available support, designing and implementing extensions to Linux are very straightforward. This support will be handy when TABC decides to maintain our re-designed cluster-compatible ftp, file, and mail servers. To maintain the security of the system, we will be designing our own authentication file with a special encryption and hashing algorithm.

We will be building our system on the personal computer (PC) architecture, as it is by far the cheapest hardware platform available today. Due to its widespread use throughout the world, PC hardware costs much less than proprietary hardware solutions. PC hardware is also tested by a larger market and therefore stability statistics are often more readily available. Likewise, performance statistics are also better understood and it is easier to make a cost/performance analysis to determine what kind of a system meets an

institution's needs. Also, TABC has abundance of old PCs that can adequately run our system.

The main functionalities of this system such as synchronization of the two computers and seamless transfer of services from the non-working system to the working system will be invisible to the user. The only interaction our system will have with the user are the GUI and the re-designed ftp, file serving and email services.