

1.1 Requirements Gathering

1.1.1 Interview

By conducting an interview with the different types of users at TABC, we gathered the requirements for our system directly from the stakeholder(s) of our project. From our research on how to conduct interviews, we found that people rarely give good answers to subjective questions such as, “What do you feel is the best part of the current system?” or “What would you like in a future system?” (Isaacs)

People tend to give good answers to things that are purely factual, or numerical, or a choice between two, or a maximum of three options. Based on this, we concentrated on asking such questions as: “What do you do when a particular situation arises?” This type of question leads to free-form answers which are not conducive to basic questionnaires, though they are much more demanding on the requirements gatherer.

We conducted an interview of our Sponsor, Mrs. Ceil Olivestone, and Mr. Ben Speiser (TABC Technology Coordinator). We asked Mrs. Olivestone mostly factual questions, but did intersperse some subjective questions in order to get a better ‘feel’ of TABC from her point of view. She directed us to Mr. Ben Speiser to answer some of the more technical questions.

Persons in this interview:

BS: Mr. Ben Speiser

CO: Mrs. Ceil Olivestone

SP: Sriram Polepeddi

This section is for getting the background on the current environment and particular needs of TABC, which we will use in developing our solution.

SP: How often do you use the computing infrastructure at TABC, and for what applications?

CO: *Mostly email, web browsing, most Office Files (Word, Accounting, etc) are stored locally.*

SP: Which of these applications would you describe as mission critical, or necessary for the basic functioning of the organization?

CO: *The Accounting package. I can live without email or web-browsing, or anything network related.*

SP: Which applications at TABC does the general population use most often?

CO: *Email, Office Apps, web-browsing, shared network data drives for faculty and students, and the most complex software we have is the Database.*

SP: Which of these applications would you describe as mission critical, or necessary for the basic functioning of the organization?

CO: *Email is critical as we need it to communicate with the parent body. The web is used heavily by the faculty and students. The Database is used by one School Administrator.*

SP: How would you rate the computing knowledge of: the general population at TABC, the Administration, the Faculty, and the person(s) responsible for maintaining the computing infrastructure?

CO: *Among the School Admins: mostly very little, but there are a few good-to-very knowledgeable persons. Among the Faculty, it is mixed, the younger ones are better than the older ones. The children are very advanced, only a few are not. And I would rate our Tech Department as being very very high on the computing knowledge scale.*

SP: How often do you receive complaints with the Network?

CO: *Not so much about the network, but more about the hardware, esp. the older hardware. But Ben Speiser would know more.*

BS: *It ranges. Sometimes none, sometimes as many as eight complaints per week . It depends mostly on when the administration or faculty have particular deadlines approaching, and things just don't seem to be operating under the strain.*

SP: What is the procedure step-by-step that you follow from the point of receiving the complaint to its final resolution?

CO: *All Ben, Ask him.*

BS: *Well, I first try to determine the severity of the problem. If it's no more than a nuisance, I try to explain how it's low on the priorities and I explain a current workaround. If it is important but not immediate, it goes onto my "to-do" list and gets taken care of during the next round of "fix-it" time. If it is of high priority or can be easily resolved on the spot, it is taken care of immediately if time allows. The problems are from too great a range to elaborate anything beyond this damage assessment phase, though.*

SP: If a disaster struck your current setup that prevented it from providing these services, how much would you lose in tangible terms?

CO: *Cannot put dollar terms on it, but maybe in terms of repair costs.*

SP: And in what other intangible ways would TABC be affected by downtime?

CO: *We can function as a school, as the education, other than the computing department, is not that computing based. But administratively, it would affect us. We would not be able to access the kids' schedules, teachers' contracts, billing info, etc. Ben would be able to provide even more info.*

BS: *Every single thing I should be doing would get put on hold. I would not be able to teach classes, make phone calls, grade tests, call parents, monitor students in the computer lab; deal with students who break school computing policy... the list goes on and on. I remember when our old ISDN router failed because our ISDN line was hacked - our entire DHCP system shut down - I did nothing for two days other than give PCs static IP addresses. Afterwards, I had to switch them back after our DHCP came back online.*

SP: Have you ever had a major issue at TABC, which necessitated an expensive repair solution?

CO: *At least not in the three years that I've been here.*

SP: What issues do you presently have with the current infrastructure even when it is operational, if any?

CO: *Well, no specific issues, other than our low tech knowledge which does cause some issues.*

This section is related to the possible solution we will develop.

SP: Is any solution intended to add to TABC's current infrastructure or to replace existing hardware?

CO: *Ben Speiser would know that*

BS: *Currently, we want to use it to augment our infrastructure. Our current authentication servers do double duty as file servers. In some cases, this makes sense, but it also means that they can be overtaxed. We don't intend to replace any hardware or infrastructure, but we do intend to augment it.*

SP: How important, if at all, is it for our solution to NOT require additional purchases on TABCs part?

CO: *Very important, not much in the budget, but I can perhaps spare about \$100, in addition to all of the parts that are already available in our storage.*

SP: Who do you see as the people who will need to interact with the system on a daily or occasional basis?

CO: *That would be Ben Speiser.*

SP: Where in TABC are you planning on implementing this product?

CO: *Ben Speiser would know that.*

BS: *First off, in the publications room, as a standalone server. Ultimately, I would like to be able to deploy it as a cluster in both the main office and the computer lab.*

SP: Would you prefer a system that gives decent performance on all hardware, such as the old Pentiums in your Hardware Inventory, or a system that can give excellent performance on a smaller subset of your Hardware, such as the Pentium IIIs? We would likely still be able to use whatever hardware you already have, without requiring the purchase of additional equipment.

CO: *Why would I want a slow application if I can have an application that would much faster on better hardware.*

SP: So would you say that you are more concerned about maximizing the productivity of the system than simply maximizing the number of systems that this application could be run on?

CO: Yes

SP: Who at TABC would have the final responsibility for the management of the system? i.e. Is it important for the Administration at TABC to be able to control this application if need be, with their own special privileges, or would the entire administration of this system be left to the person(s) currently maintaining it?

CO: *That would be Mr. Speiser*

SP: Would you or the administration need to be appraised of any and all actions that the system performs or would that again be something left to the administrator of the system?

CO: *Yes, though the administration would not be doing any of the technical activities related with it, we would still need to know what is going on with it.*

SP: How soon would you need this project to be implemented?

CO: *Though, it would be good to implement it this semester, this is not a priority*

SP: What are your expectations of the project once it is completed?

CO: *Well, I do require proper documentation so that anyone who looks at it can administer the system.*

CO: *Other than that, what would I do in case something happens to it 3 years from now?*

SP: *We are building upon open source components that have been widely tested by an entire legion of Linux programmers, so they are very stable. The documentation we give you will allow anyone with even a slight knowledge of Linux Administration to rectify most problems. If this doesn't work, there are plenty of Linux Tech Support Professionals who could take care of any problem.*

CO: *Ok good.*

SP: *Would you have any need for the source code of this project once it is implemented?*

CO: *What can I do with that?*

SP: *Well, it would allow you to make low-level changes in case you wanted to tweak the system even further. But this would assume that you had a Unix programming expert at your school ready to do this.*

CO: *Well in that case I don't see the need for it.*

SP: *Would you want the source code for your computer science classes for your students to work on?*

CO: *I don't know I will ask Mr. Speiser about that and let you know.*

(I later asked Ben, and here is his response.)

BS: *It might be nice to use some of the code in my class, but this is something I don't think I would want students tampering with. I am very security-conscious, and knowing that a student would need to be trusted with the underlying code that makes the server accessible is concerning. In the long run, I would probably prefer (from a security standpoint) that the software source code be unavailable if possible.*

SP: Thank you Ceil. I appreciate the time you have given us.

1.1.2 Observation

Torah Academy of Bergen County (TABC) is a private school based in Teaneck, NJ. The school teaches classes from grade 9 to 12. Students here undertake very challenging projects under guidance from top-notch faculty. There is a science lab at TABC for working on projects involving lasers, jet spectroscopy, monochrometers, telescope positioning systems, etc. Everyone at TABC uses the computing infrastructure in some form or the other. The main users we noticed were the school administration, faculty, and students. Though they all utilized different applications, they all had the same view of the network; i.e. they all saw their applications as files kept somewhere on the network. The only difference between them was the permissions they had to access the network. These permissions were granted to the various users through authentication servers kept in the publications room, the computer lab, and main office.

While students use the latest available PCs for their work, many important administrative responsibilities are handled by older PC architectures. This poses a threat to the TABC's data possibly hurting their operations in case any problem arises with these machines. They are all aging, some are cobbled together using parts harvested from other computers, and none use backup hardware such as tape or CD-ROM. There are at least three servers that are in dire need of a system reliability solution. The first is the office server used by the entire administrative staff to store daily data files, Peachtree Accounting Spreadsheets, etc. The second is a server used to house TABC's newsletter and yearbook work. The third is a server that houses the shared network drives of each faculty,

administrator, and student at TABC. Also, there are servers that perform vital functions for the school, such as the firewall and print servers. These will also use a system reliability solution if it is cost-effective and fully tested.

1.1.3 Use-Case Scenario Diagram

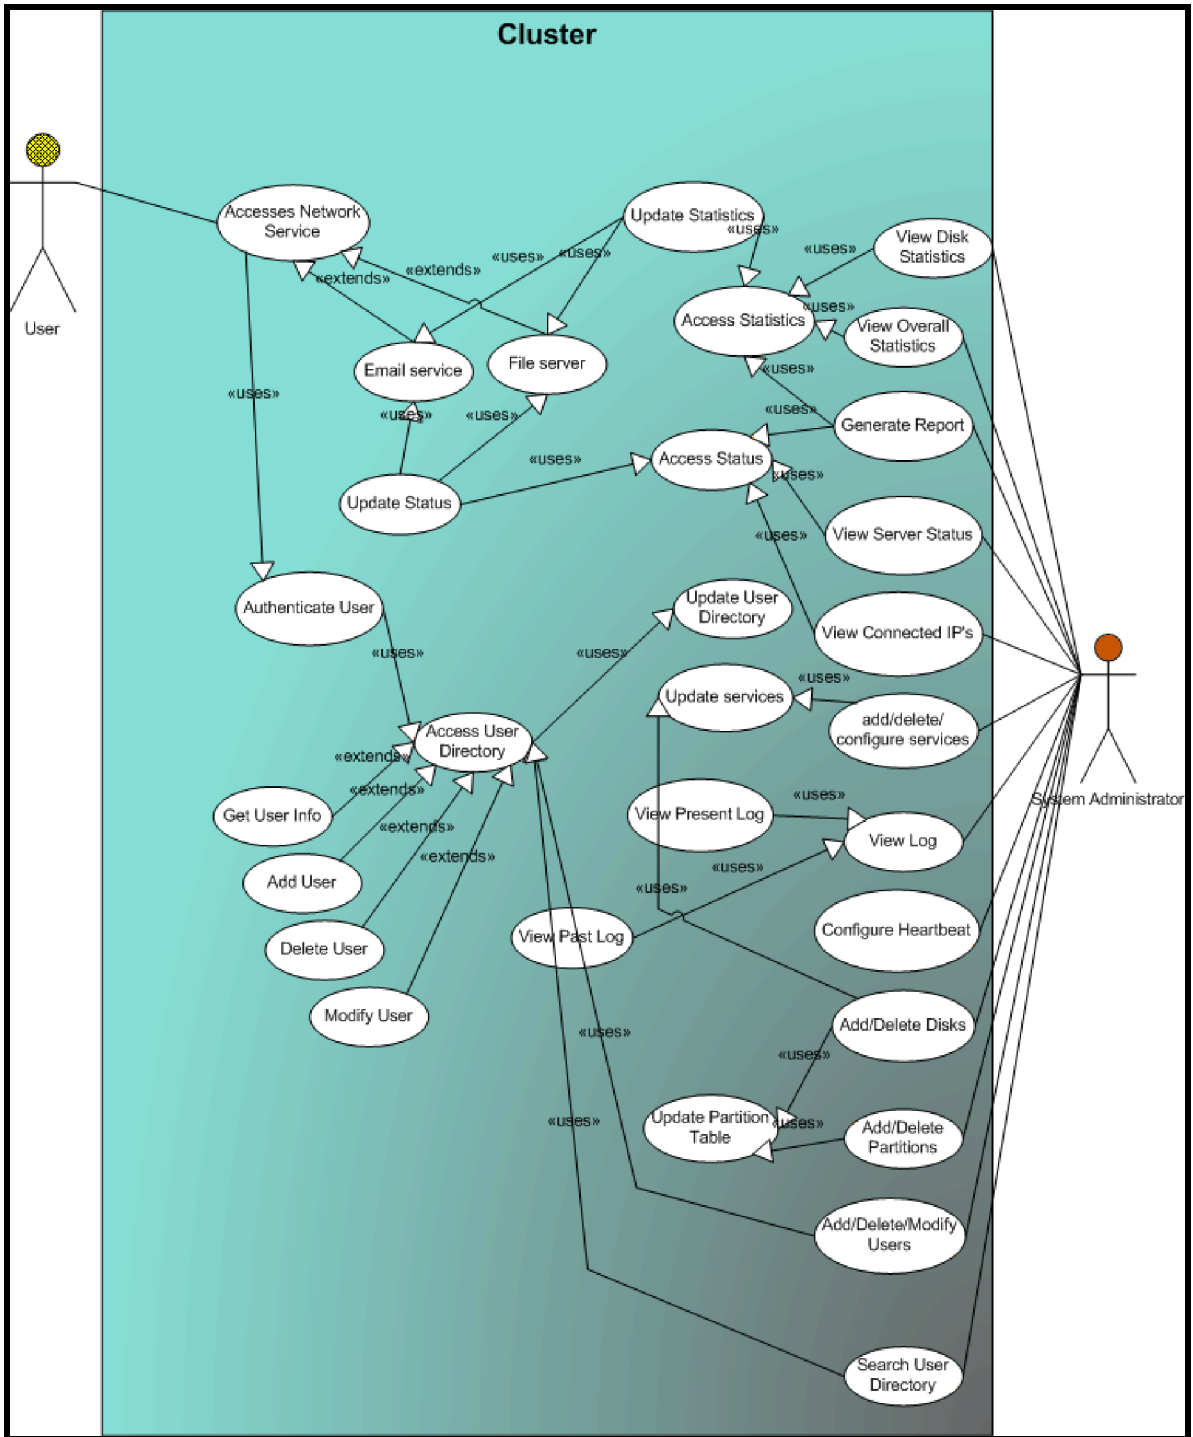


Figure 1: Use case scenario

1.1.4 Brainstorming

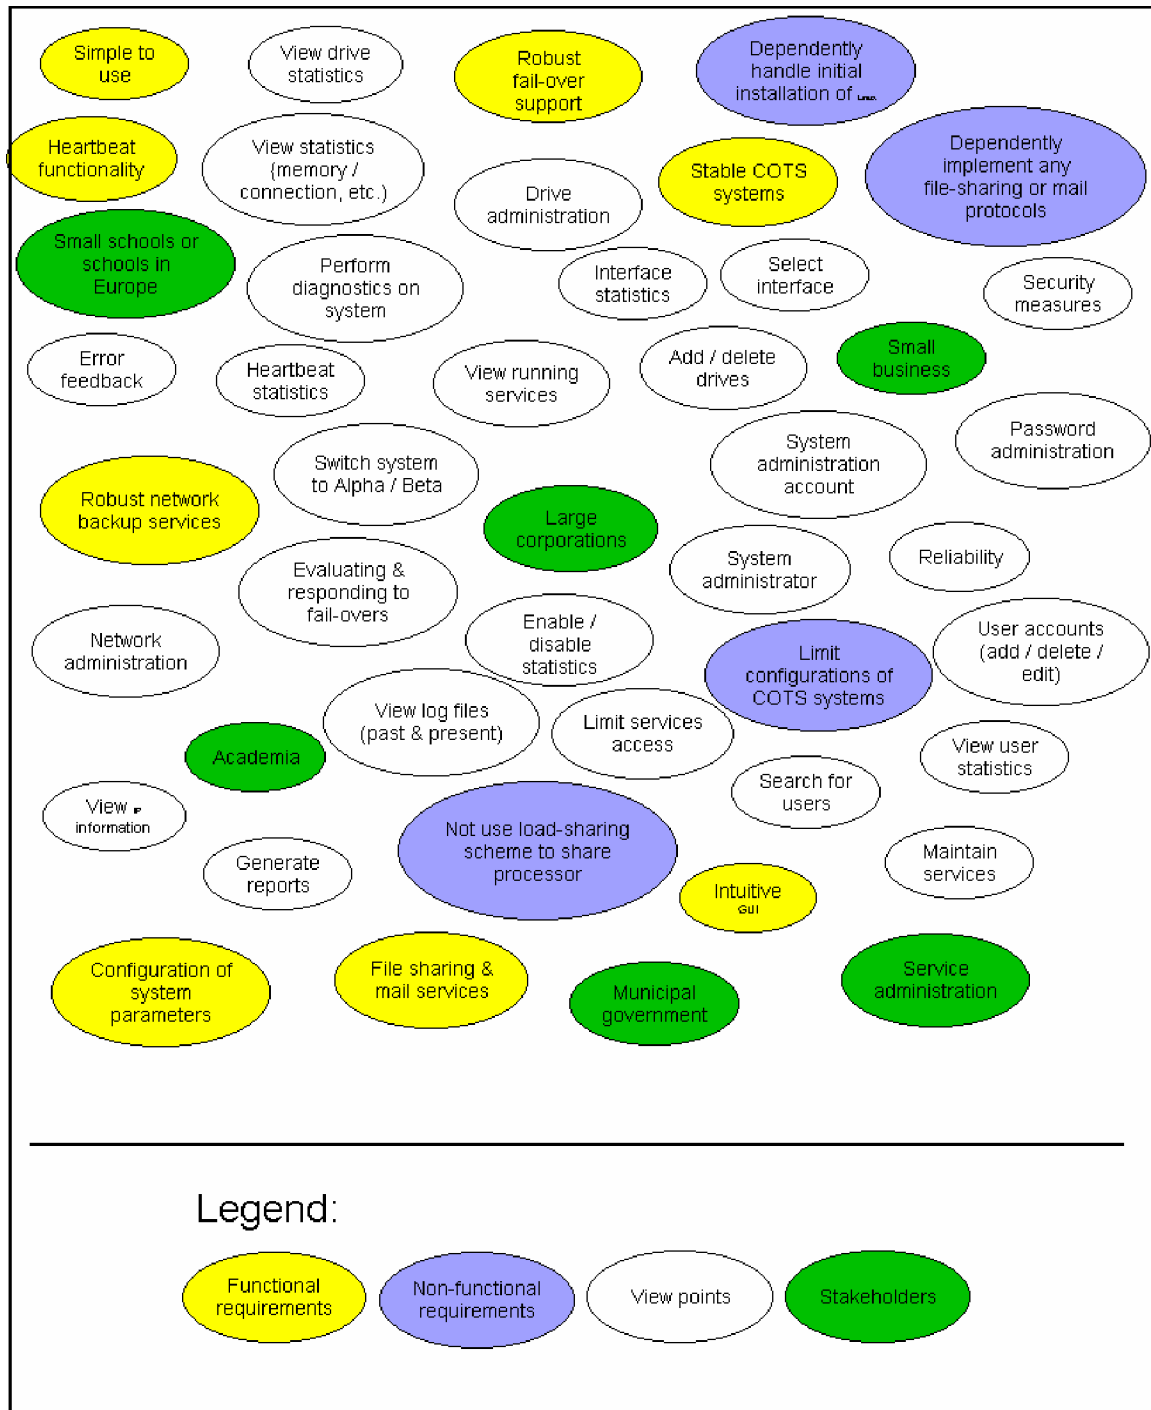


Figure 2: Brainstorming

1.1.4.1 Brainstorming Prioritization

Each user was given 100 points to distribute between these 4 major categories within the project. As the project progresses through the software development life cycle, more efforts will be focused on the categories gathering the highest scores. The reason for this is the user community has voted on what they believe are the more important aspects to a successful project.

User	Security	Reliability	Easy to use	Reporting abilities
Large Corporate User (William)	20	50	20	10
Small School User (Ben)	20	40	25	25
Small Business User (Ram)	20	35	30	15
Academia User (Bayo)	15	30	15	40
Municipal government User (Mehul)	40	35	20	5
System Administrator (Ryan)	10	50	0	40
Totals	125	240	110	135

Table 1: Brainstorming prioritization

1.2 Post Requirements Gathering Problem Statement

The main concept has not changed much from the earlier problem statement as the goal of designing a clustering system that works as a simple and reliable backup solution still exists. All of the features presented here build on the previous problem statement and do not replace its features in any way.

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. And are providing a graphical user interface (GUI) for the person administering the cluster. 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. 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, and after consultation with one of the stakeholders feel that we should optimize the code for a Pentium III Architecture. Also, the administrator at TABC wishes to be alerted by email in case of a problem with the cluster. And in addition to the basic services such ftp, file serving, and email, we have found a need for specially tailoring DNS for the cluster as well.