

i.d. miller  
6114 La Salle Ave Box 543  
Oakland, CA 94611  
<http://idmiller.com>

**This is a DRAFT**

**Implementing a Psychological Survey using Open Source Software.**

[http://idmiller.com/media/publications/oss\\_survey\(2006\).pdf](http://idmiller.com/media/publications/oss_survey(2006).pdf)

© 2007 Ian Dennis Miller

Abstract: A variety of Open Source products were combined to provide an online survey function for research in social psychology. The specific technologies, the way in which they were combined, and the challenges doing this are discussed, as well as the requirements that were fulfilled.

## Introduction

I was hired in July of 2003 to program an online survey for a psychological experiment being conducted at the University of California at Berkeley by Elizabeth Page-Gould, Professor Rodolfo Mendoza-Denton, and Professor Linda R. Tropp. The timeline for the project was fairly compressed and inflexible, and the available resources were scarce. Taken together, these interesting constraints influenced the entire direction of the work towards Open Source Software because:

- Open Source Software (OSS) is free on its surface; OSS is a tradeoff between budgeting for labor and budgeting for commercial software
- commercial software is more difficult to modify internally (it is "closed source")
- commercial software is generally geared towards projects with high-performance requirements, which was overkill in this case, to the extent that the timeline might actually suffer for it
- I was personally familiar with the OSS offerings because my undergraduate education favored OSS as a learning tool, because of the visibility of the source code

This paper is organized into two sections: the features required by the project, and the technologies used to implement those features. The technologies have since improved, but the features that are fundamental to any survey are orthogonal to the technologies that might implement them. In short, pencil and paper (old technology) will still suffice, and can even benefit from the discussion of survey features, but the technology currently available gently suggests that a better approach exists.

## Technologies

For the purposes of this discussion, technologies are both hardware or software. From a planning perspective, a unit of technology is a "black box" whose function is understood without necessarily understanding how it functions. The automobile is a prototypical black box technology: I expect the car to stop when I hit the brake, but I don't worry about how the anti-lock braking system detects and reacts to road conditions.

Planning the software architecture of an on-line survey follows the pattern of many web applications, which is frequently called Model-View-Controller (MVC). The MVC pattern, applied to web applications, generally calls for the following technologies: on the client computer is a web browser, and on the server computer is a web server, a database server, and the application that controls it all.

The term LAPP Platform (or LAMP) is commonly used in the Open Source web application literature due to its relation to the MVC pattern, where LAPP stands for Linux, Apache, PostgreSQL, and Perl. The LAPP platform of software technologies, which will be described in detail, provide a web server, database server, and application programming language, all of which are open source. All of these technologies are built on that most-general-of-all technology: the Intel x86 processor and the "generic PC."

In discussing the technologies involved in this on-line survey project, I will work from lowest level to highest, providing a peek at the insides of these "black box" technologies as they related to the on-line survey project.

## Hardware

I used two physical computers in this project: a development machine, which I administered, and a production server, which was controlled by the lab that performed the research. Hardware is still an issue for any lab or project, and it is possible to create higher demands than your hardware can sustain.

Unfortunately, it is impossible to provide any rule-of-thumb in estimating hardware requirements because it is depressingly easy to code inefficient web applications.

Consider the following requirement: up to 8 users will simultaneously access a database that will accept approximately one page of data per minute from each user, or 8 pages per minute. This means that, on average, the survey must be able to sustain one page every 7.5 seconds. For perspective, Amazon manages tens of thousands of pages every 7.5 seconds, and this 8-person survey only has to handle one page in that period of time.

In terms of hardware, of course, Amazon has thousands of computers, and you might only have a single, generic PC from 1997. I'll go out on a limb and conservatively estimate that the 1997 PC probably won't suffice for our 8-user requirement, as meager a requirement as it might be. Also, keep in mind that decade-old hardware suffers from other aging problems, too.

All programming ineptitude aside, I will venture some rule-of-thumb estimates (anecdotally taken from personal experience) for the 8-user survey:

- newer is better: Intel P2 400MHz is scraping by, Intel P3 800MHz is comfortable, Intel P4 2.5GHz is quite comfortable.
- no matter the age of the computer, maximize RAM: 128MB is scraping by, 256MB will be adequate, and 512MB will be comfortable.

In practice, my hardware consisted of two machines for implementing this survey. Here are the exact specifications:

i.d. miller "development" machine:

- Operating System: GNU/Linux
- Processor: Intel Xeon 1.8GHz P4 (2 processors)
- RAM: 1GB (512MB, 2 DIMMS)
- Hard Drive: 36GB SCSI 160

UC Berkeley "production" machine:

- Operating System: Windows XP SP2
- Processor: Intel P4 2.53GHz (1 processor)
- RAM: 512MB (1 DIMM)
- Hard Drive: 60GB ATA 100

The hardware was more than adequate for the needs of the survey. However, such a setup had interesting consequences, not the least of which was the difference in operating system, which I will discuss in greater detail.

## **GNU/Linux**

By now, the Linux kernel may need no introduction, as it appears in the media with increasing frequency. Linux, or more specifically, GNU/Linux, is the combination of the Linux kernel with a wide range of tools that approximate a Unix environment. The kernel, as well as the GNU tools that make it usable, are all Open Source products. The nature of the Linux kernel is to serve as an interface to a computer's hardware, and in this case, the less exotic the hardware, the better.

Hardware compatibility issues did pop up occasionally on the development server, although this did not affect the project timeline. To illustrate, the Linux kernel (in 2003) did not support certain error

codes that were generated by the processors in the development machine. Because the hard drive was physically installed in such a way that it restricted airflow to the processors, the processor temperature exceeded its threshold, and generated a “Machine Check Exception” that the Linux kernel was unable to handle, which resulted in a very nasty, sudden machine shutdown. The solution to this problem was twofold: moving the hard drive to restore airflow, and waiting for the gradual update of the Linux Kernel, which can now respond to temperature events by lowering the processor speed.

The common criticism of GNU/Linux comes from the perspective of the desktop user experience. This is largely irrelevant in the case of online web surveys. In fact, considerable benefits come from not running a graphical desktop at all, including better security, more predictable operation, and better utilization of memory. The desktop argument aside, GNU/Linux supports a suite of common server tools for making web applications: Apache, Perl, and Postgres, all of which are discussed in detail.

GNU/Linux is an excellent operating system, but as the anecdote warns, unexpected interactions and fringe operating conditions can result in serious troubleshooting. In my experience, this is easily avoided by using mainstream, modern hardware. GNU/Linux performs well as a dedicated server operating system, and is a good platform for building web applications.

When people talk about Linux, they often discuss the wide variety of Linux “distributions”, which are also referred to as flavors. To briefly comment on Linux distributions, some are tailored to new users and others to experts, so it is worth it to research before making a decision. During this project, I used Gentoo Linux, which is only appropriate for advanced users, but I currently use Ubuntu Linux, which is much more accessible. I would recommend Ubuntu to OS X users (the underlying operating system is quite similar), and I would even venture that the switch for Windows XP users can be managed.

Back in the real world, labs can't often afford a technician who knows Linux, and may not be lucky enough to have a volunteer with functional Linux administration skills, either. Although practically all computers can run Linux, a measurable expense manifests through the *support* of Linux. If “computer support” isn't literally in the budget, it implicitly appears under other headings as “wasted time.”

## Windows XP

Back in the real world, labs can't often afford a technician who knows Windows XP, and may not be lucky enough to have a volunteer with functional Windows XP administration skills, either. Although practically all computers can run Windows XP, a measurable expense manifests through the *support* of Windows XP. If “computer support” isn't literally in the budget, it implicitly appears under other headings as “wasted time.”

In contrasting these two statements about “wasted time,” the difference is clear: computers are difficult to use, and that value can be measured in terms of support time. I assert that support time is partially a function of troubleshooting knowledge and training. Windows XP (often called XP for short) is a monster to learn, but it so happens that a lot of people in the academic world are reasonably well-trained at XP. As a desktop operating system, it is easier to use than most alternatives. Therefore, Windows XP is probably the most common operating system used in academic labs, and that is the kind of computer you can expect to find lying around.

In the case of the on-line survey project, one requirement was that the survey be hosted on a machine running XP. This requirement is not too difficult to derive: for the purposes of data security, the machine must be kept physically secure, and the lab already had a locked room to store it in. Furthermore, without budgeting for a dedicated machine, an existing machine would have to be reused, and where else would that machine be but in the lab? Unfortunately, from the Unix perspective, XP is not well suited to be a server operating system.

## Cygwin

My preference is to develop for GNU/Linux, but the project required Windows XP. I chose to bridge this gap using Cygwin, which is a Windows XP tool that provides a similar experience to Linux. Cygwin also serves as a gateway into a world of pre-compiled Linux applications that run within the Cygwin environment. To relate this back to the LAPP platform, it is possible to swap out Linux for XP, because Apache, Postgres, and Perl are all quite functional under XP with Cygwin.

This level of functionality is provided for by competing Open Source projects, too, such as Indigo Perl or XAMP. In the end, there is only one requirement: that the computer runs Apache, Postgres, and Perl. The Operating System is, in a sense, secondary. However, committing to XP and Cygwin has its consequences, both positive and negative.

On the negative, this means you must monitor yet another mailing list to watch for Cygwin security updates, and you must act on them. On the positive side, though, Cygwin really behaves like Linux in many ways, so remote system administration becomes possible (via OpenSSH). Over the course of the survey's lifespan, the need to remotely administer the web server will inevitably arise, and fixing it faster is always better. It is a different process to remotely administer Indigo Perl and XAMP, and I personally feel less comfortable with it.

A major difference between XP and Linux that Cygwin underscores is the process for automatically launching an application at bootup. Also, Cygwin appears less like Linux the deeper you dig, so depending on the requirements of your web application, you may find major architectural differences that are unexpectedly painful to code around.

For most on-line surveys, this won't be a problem, but consider the following: as part of a data encryption scheme, I proposed using a certain algorithm that depended upon a specific library of code containing exotic math functions. While I could easily install these math functions on Linux, I ran into deep architectural difficulty with Cygwin and XP. In short, a different encryption approach was required because Cygwin wasn't similar enough.

## Apache web server

The Apache web server (often called just Apache) is a special Open Source project. Beginning as just a web server, through the explosion of the Internet, the project has blossomed into the Apache Foundation, which is an incubator for dozens of other projects now, including significant technologies like the Apache Tomcat server. The Apache Foundation is very alive as a community, and the Apache web server appears to have a solid future ahead of it.

In building a web application, the web server manages the connection between a client (who is generally using a web browser) and the web application (which, in this case, is a survey program). Apache performs admirably in its web server capacity by providing accelerated interfaces for many programming languages, including PHP and Perl. In communicating with clients (including potentially malicious clients), Apache has historically been quite secure, with infrequent but timely security updates.

As with any web server, security will be a major issue. Principally, a web server is software *running on your machine* that allows a remote user to control it. This introduces all sorts of complications, and it becomes extremely important to limit the actions a client can perform. Apart from your Hollywood-style "hacker" scenario, in academic research, a potentially serious problem is exposing sensitive data. Specifically, the layers of interaction in any moderately complex web application potentially introduce new avenues for an information leak, which means, for example, that a completely secure web application can be undermined by an insecure web server configuration.

The web server is among the most visible security components in an on-line survey, but it is by no

means the end of the story on security. In the same way that a secure application is undermined by an insecure web browser, a secure web browser is undermined by an insecure operating system. As far as the OS is concerned, Apache installs all flavors of Linux (or comes pre-installed). In the case of XP, it is a mixed bag, but as was discussed before, Cygwin, XAMP, and Indigo Perl all present options.

## PostgreSQL

PostgreSQL is an opensource database server. While it faces fierce competition from MySQL, a peer in the Open Source database arena, I have backed PostgreSQL for several years due to a single advantage it alone possesses: ACID compliance (which stands for Atomicity, Consistency, Isolation, and Durability). Whereas MySQL can fall into a corrupted state, to be recovered later by their recovery software, PostgreSQL will never enter such a state in the first place. To accomplish this, PostgreSQL has historically showed slower performance than MySQL, which is an important factor in the commercial sphere.

However, because the performance requirements were fairly modest, and the cost of losing any data was immense, the decision to pick PostgreSQL over MySQL seemed obvious. PostgreSQL installs easily enough using Cygwin, and is generally simple to install under Linux, as well.

A number of unique factors affect the database aspect of an on-line survey. A database server supports a finite number of simultaneous connections, but the behavior of Apache is to hang on to connections longer than is necessary, which can unnecessarily impact performance. It is also a non-trivial task to design a functional database table for the needs of a given on-line survey.

## Perl

Perl is a programming language. Syntactically, it is pretty similar to C++. In execution, Perl is an interpreted language (as opposed to compiled). It is modern, in the sense that it supports object orientation (ugly, but true) and it provides memory management and garbage collection. Finally, Perl is very social, perhaps unlike anything before it. Through CPAN (the Comprehensive Perl Archive Network), a massive library of code is available for reuse. Although it is at first difficult to reuse other people's code, the payoff is massive.

Returning to the LAPP platform, the final P stands for Perl. The acronym is clever, but in practice, Perl can be swapped out for PHP, Python, Ruby, and any number of other Open Source programming languages. Each language has its fans and detractors, although in the web sphere, there is particularly high use of PHP. Personally, I prefer to think of an MVC application in Perl terms, although there have been projects to adapt MVC to PHP, for example.

As languages go, Perl has a notorious reputation for being ugly and illegible. This is easily true, and just as easily false. From experience, I know that it is possible to elaborate byzantine structures and labyrinthine tangles of code, complex enough that they work but you have no idea how. Also, from experience, I know that when your code is that complex, you ultimately have a poorly organized understanding of the underlying concepts. I will confess that the on-line web survey suffered from a touch of the labyrinth and deviated rather far from the MVC pattern. These often go hand-in-hand.

Perl interacts very well with Apache and with PostgreSQL. In the case of Apache, the `mod_perl` project provides the ability to directly integrate Perl code into Apache, resulting in a massive performance payoff. PostgreSQL is accessed from Perl through the DBI module, where Perl DBI can improve performance again through `mod_perl` and Apache. In short, Perl is an excellent language for coordinating web server and database server activities because it integrates deeply with both.

## Web Browsers

The web browser is generally the target of a web survey, or rather, it's the person who is using the web browser. Part of the challenge is in ensuring that nothing interferes with this process, so that everyone sees the same fundamental thing. Of course, this is notoriously difficult.

Much of the difficulty started with the so-called "browser wars" of Netscape Navigator versus Microsoft Internet Explorer. An international standards body, the W3C (World Wide Web Consortium), established a series of specifications for HTML (Hyper-Text Markup Language), upon which the so-called World Wide Web has been built. Microsoft and Netscape both violated the standards, creating a rift in the HTML language, and an endless nightmare of compatibility problems for web designers as a whole, but for on-line survey engineers specifically.

Adding to this mess is confusion about what, exactly, a web browser is capable of. To answer the first question that comes up, millisecond-scale reaction times are not possible to record using a stock browser. Experiments that purport to record millisecond-scale reaction times must rely on Java applets, which can be embedded in the browser, but which still cannot reliably guarantee such accuracy.

Although it is difficult to consistently display web pages properly, and reaction times cannot be recorded for an on-line survey, it is still possible to use the web browser to collect certain, meaningful kinds of data.

Web browsers also play a key role in encrypting communication. By serving as one endpoint in an SSL-encrypted channel (with the Apache server as the other endpoint), the web browser helps to minimize the chances of third-party eavesdropping.

## Features

The vocabulary varies on this topic, ranging from "requirements" to "features" to useful frameworks such as "Use Cases", but for the purposes of this paper, "Features" are the things that the software accomplishes. The primary feature of a minimally functional survey is that of data collection; participants enter data, and the survey records it. For consistency in vocabulary with other work, "participants" are the same as the deprecated term, "subjects."

## Data Tables and Recording Data

To accomplish the survey's primary objective, participants responses were recorded in a database. There are two primary ways to accomplish this: in the manner of a "spreadsheet", or in the manner of a "journal". A spreadsheet would contain a column for each datum to be recorded and a row for each participant, and each data point that is recorded is entered at the appropriate location in this matrix. A journal, on the other hand, is a running list of each data point entered, and it contains only four columns: the identifier, a timestamp, the variable name, and the recorded value of that variable.

For this survey, a journal approach was used.

In future surveys, both approaches were used.

## User Identification

## Security and Privacy

## Dynamic Elements

Dynamic Elements are elements appearing in the survey that are determined at run-time, or "on the fly." This may be driven by the condition the participant is assigned to (randomly or otherwise), or by answers the participant previously provided. In the case of this survey, both were the case: survey condition was determined by the background of the participant, and the names of participants' friends were reused a number of times throughout.

An example of a dynamic element is to create a survey item that contains a "placeholder," which is blank until actually viewed by the participant. Upon viewing, the placeholder is replaced with, for example, a word representing their ethnicity, as in:

"Imagine you are out tea with a \_\_\_\_\_ friend"

which becomes, depending on the condition:

"Imagine you are out to tea with a caucasian friend" (though this can just as easily be "latino/latina" or any other scientifically indicated phrase).

Although it might appear trivial to accomplish this, consider the following example, where a participant has previously indicated that their friend is named Alfred:

"How close do you feel to \_\_\_\_\_?"

which becomes:

"How close do you feel to Alfred?"

In this case, the friend's name has been recorded in the database, so it must be looked up in the database in order to be displayed. Depending on the volume of concurrent users that the survey is actively communicating with at that moment, the total number of entries in the database, and the design of the database, this lookup to retrieve the name can take a non-negligible amount of time.

It is a matter of user experience that this retrieval time be minimized, because in the extreme, this could potentially alter the experience of the participant in an experimentally noisy manner. As the number of users increases, the delay might increase, causing participants who take the survey later to experience a longer delay than those who took it earlier.

The solution is to properly design the database with this in mind. Because this was an issue of this survey's table design, a feature of Postgres called an "index" was used, which has the effect of pre-sorting the data tables based on a particular field. Because users were uniquely identified by an integer, the table was kept in a pre-sorted state on the basis of this identifying integer.

To take advantage of this, database entries were first filtered by user identifier, which had the consequence of capping the total number of fields to be scanned for the friend's name. This was verified by including the "render time" of a page in special administrative debugging output, where render time is the amount of time it takes the page to be displayed.



## **Loops and Forks/Skip Lists**

## **Multiple Sessions**

## **Obtaining Access**

## **"Scales" versus "Pages"**

## **Appearance: HTML and CSS**

## **Acknowledgments**

This survey work was not conducted in a vacuum, and since 2003, I have participated in several projects which have extended the work I began here. The following people have been involved, either by hiring me to work on their projects, or in the case of Guido and Aaron, to collaborate on a commercial architecture for academic surveys.

- Liz, Rudy
- Diane, Jonathan, and CSHE
- Guido, Aaron