

Sahana Phase II: Master Document

Contributors: Mr Chamindra de Silva, Dr Sanjiva Weerawarna, Mr Jivaka Weeratunga

Version 0.5

DRAFT – Not to be published

Table of Contents

Sahana Phase II: Master Document.....	1
The Vision for an Ideal Disaster Management System.....	1
Why Open Source?	2
Current Sahana Project Status.....	3
The Future Direction of Sahana.....	4
Sahana Phase II Components.....	5
The Potential for Humanitarian FOSS.....	5
Scope, Assumptions and Constraints on System.....	6
High Level Work Breakdown and Roadmap.....	6
Getting Requirements from NGOs and Gov	7
ASP and Voice portal extensions to Sahana system.....	7
Domain and Architecture Leadership Team and relevant roles:.....	8
References:.....	8

Introduction

The Tsunami that hit Sri Lanka on December 26th resulted in a massive outpouring of support for the relief of the nearly one million people that have been affected by it. When literally thousands of people from every conceivable multilateral organization and from many other places arrived here to help, it became clear immediately that without information technology it would be impossible to coordinate their efforts to maximize the impact on the affected people. The Sahana project was born.

Despite the tremendous value this type of software can bring to disaster management, there are only very few systems that exist today and none of them are widely deployed. In fact, the most widely used system appears to be non-Web based and using completely out-dated technology. While there are indeed various specialized parts that exist, there does not exist a single cohesive system that organizations such as United Nations Disaster Assistance and Coordination (UNDAC) deploys at every disaster situation they go to.

The long term objectives of Sahana are to grow into a complete disaster management system, including functionality for mitigation, preparation, relief and recovery.


The Vision for an Ideal Disaster Management System

Imagine if you were in charge of managing the information flows in such a humanitarian disaster, which is critical to an successful relief response. Also imagining this has happened in a developing nation, with a low (or damaged) IT/Telecom infrastructure. Such was the situation faced by much of the areas affected by the recent Asian Tsunami. What would you look for in the ideal system?

Well for one it should be quickly deployable using commodity hardware and easily accessible (preferably free) software. It should be quickly customizable to the particular disaster.

The system should be easily accessible to all coordinating and relief providing organizations (Government and NGOs). Preferably they should have interfaces customized to their specific needs though everyone should be working off one data set to avoid information redundancy. Customization would include localization into the indigenous languages and having voice portals for counties with low literacy rates. The interfaces will include metric oriented reports customized to the decision makers of the various organizations and also provide transparency to the aid sponsors to give a degree of comfort on effectiveness of their sponsored organization. At the same time the information should be secured and should reduce the possibility of data privacy violation or abuse for adverse reasons (e.g. identify theft, property theft).

The system should operated in the field though portable devices (e.g. PDAs, mobile phones, satellite phone) and be able to work in a disconnected model as mobile stations where access the central data is not always available.



Though we have just covered emergency relief here, the scope and benefit of a system is wider to the areas of recovery and rehabilitation

Why Open Source?

Very few countries and organizations today can afford to invest a lot of resources in disaster management when there is no disaster present. While this is obviously true of poor, developing nations, it is also true of richer, developed countries as well because there are always higher priority items that need the funding. Worse yet, even if there are some national scale systems that may get deployed, it is very unlikely that regional and local level systems will ever get deployed if they cost any significant amount of resources. Besides it almost seems unethical to restrict licenses to such software at a point of a humanitarian disaster, so there see very little commercial motivation to build it. This is what we see in the world today – while disaster management software is critically needed, there is no complete commercial

or non-commercial software solution that is widely available.

Going the open source way can address both these concerns. Using the open source development model, it is possible to develop this software at a much reduced cost compared to pure commercial development models. This is true because while commercial entities are not willing to invest into these systems, there are hundreds and thousands of well-meaning IT professionals who are very happy to donate a few hours of effort to helping build such systems. We are already seeing this with the nascent Sahana project. Thus if there was a small team which was driving such a project, then it is possible to get a lot of assistance from the global IT community to make those systems truly exceptional.

Going with open source approaches can also greatly reduce the deployment cost of this software in peace (i.e., non-disaster) times. The Sahana system, for example, can be deployed on any PC with just a Linux LiveCD (that is, a CD from which the entire system can be booted up and brought on-line). Thus, not only is it possible to run this on commodity, inexpensive hardware, it is in fact possible to not even have dedicated hardware around – just take any office PC and make that the “disaster management center”!

Thus, disaster management software seems to find a natural home as an Open Source solution.

Current Sahana Project Status

The current Sahana Project was quickly built over a 2-3 week period around the time of the Asian Tsunami by a group of volunteers to help coordinate the relief effort in the for Sri Lanka. It was deployed and run as a Sourceforge project (a well know open source software development portal) and about 40+ people contributed to it's development from various groups and companies. An implementation of Sahana was authorized and deployed by CNO (The main government body in Sri Lanka coordinating the relief effort) to help coordinate all the data being captured and at the end of their tenure it had captured data on 32,000+ families.

The current Sahana system exists as a collection of interconnected, yet independently usable, subsystems that interact with each other via a set of shared databases. The separate concerns and related components of the Sahana system implemented as of now are given below:

<i>Component</i>	<i>Description</i>
Organization registry	Keeps track of all organizations and the role + ownership areas they have in the relief effort.
Request management system	Database of all requests for support from various locations (camps, hospitals, etc), as well as offers of support from relief providers.
Camp registry	Registers all temporary camps, hospitals and locations setup to house the victims of the disaster
People registry	Database of missing, displaced, dead, orphans etc. (including pictures, finger prints, DNA samples) with advanced search capabilities.
Assistance management system	Database of all pledges of assistance and attempts to match it to the requests. Also records where the assistance was provided.

Component	Description
Key contacts database	A collection of key contacts for critical areas during a disaster.

The technology stack used to build Sahana was based on the well known Open Source solution stack, LAMP. This included debian GNU/Linux for the OS, Apache for the web server, MySQL for the database, and PHP/Java for the Web Application, which makes it pretty much free and open source end-to-end.

As the hardware resource requirements of this solution stack is low it was initial deployed on a standard desktop machine, but subsequently migrated to it own dedicated server. We also tested and found that it could be deployed on a resource limited PDA such as the iPAQ with OPIE (a linux distro supporting iPAQs) within 64 Mb of RAM.

The Future Direction of Sahana

Though a lot of the volunteer effort was well intentioned building a system at the point of the disaster meant we were playing a perpetual game of catchup as the relief requirements were all urgent. This meant not all components of the system was as effective as some came in too late to be utilized or the data had already been captured in other forms. Additionally the system was build or “hacked” so quickly that its reuability and extensibility were often compromised for the sack of meeting the deadline. The lesson learnt here is that we need to build these systems beforehand in anticipation of another disaster and it should be generic enough and quickly deployable/customizable to handle any disaster event. With this and all the invaluable insights into the requirement of a disaster management system the Shahana leadership team decided to work on phase to with the following driving principles:

- Bring about efficiencies in the response effort by sharing disaster related data across all respective organizations that producers and consumers of it
- Help maximize the use of funds to the end-victims, reducing the costs of coordination and tracking
- Reduce data redundancy, chaos and resulting wastage by consolidating all data on the disaster under one system and provide accessible interfaces targeted to different organization needs
- Protect the privacy of data and help reduce the opportunity of data abuse
- Provide timely summary and trend oriented reports to the principle decision makers to help them better plan the even distribution of aid and supplies
- Provide a very accessible system customized to the respective organizational needs, the indigenous languages of the country affected and even voice portals for the illiterate
- Reduce the risk of pilferage of aid and supplies by bring about transparencies on reporting and tracking, whilst at the same time easing the process of recording information
- Provide timely data and interfaces to data entry in the field where connectivity to the central system can sometimes be intermittent (e.g once a day or once a week)
- Be very simple to install and even simpler to use; anyone familiar with the Web and with filling forms on the Web should be able to use the system

The team is looking to rebuild Sahana under these guiding principles this year. Additional components we are looking to build include the Damage database, Burial

registry, distribution logistics management and the mobile station. Across the whole system we will enforce a security framework to protect data privacy and implement localization so that the system can be accessed in any language. Finally we noted an automated data transformation/import mechanism needs to be introduced as a lot of data due to the lack of awareness often first ends up being captured on paper forms, adhoc spreadsheets or web databases. This data should not be lost, rather we should have a quick and easy mechanism to transform and import the data into the Sahana system.

Sahana Phase II Components

The separate technical concerns and related components of the Sahana system are given below for phase II:

<i>Component</i>	<i>Description</i>
Organization registry	Keeps track of all organizations and the role + ownership areas they have in the relief effort
Request management system	Database of all requests for support from various locations (camps, hospitals, etc).
Camp registry	Registers all temporary camps, hospitals and locations setup to house the victims of the disaster
People registry	Database of missing, IDPs, dead, orphans etc. (including pictures, finger prints, DNA samples) with advanced search capabilities
Assistance management system	Database of all pledges of assistance and attempts to match it to the requests. Also records where the assistance was provided
Damage database	Records all damages to property and provides an estimate of the cost of the recovery effort
Burial registry	To record burial site information for the benefit of future reference
Data consolidation module	Quickly transforms and synchronizes data from disparate sources into the central database
Reporting module	Generates custom reports and statistics as required for the different organization groups
Mobile clients and station	Allows for access to Sahana through mobile devices in the field and also allows for disconnected operation using localized databases
Admin and security module	Control and review access rights to the sensitive data into the Sahana system
Internationalization customization module	Customization module to quickly internationalize and localize system for countries and races affected

The Potential for Humanitarian FOSS

This is but one example where FOSS based solution are being used to alleviate human suffering, but this concept finds a natural home not just in disaster management, but in a superset that extends to humanitarian ICT or any other ICT requirement which concerns the improvement of human welfare. We found that the currently taxonomies of projects on well known open source repositories like sourceforge or freshmeat does not presently allow us to bucket such project easily and often get dropped into a misc classification bucket. However we believe there is a lot of potential for growth in this area, namely “Humanitarian FOSS” and if positioned and promoted well they should be many volunteers flocking to build and contribute to such projects globally, especially as the open source community operates with a strong set of ethics for the benefit of the community at large.

Scope, Assumptions and Constraints on System

While the system is quite flexible to handle different humanitarian disasters it is important to specify the boundaries of this system and assumptions made to allow us to focus.

System scope

1. Only to be used in relief, recovery and rehabilitation effort and not to be extensible beyond this
2. This system is not supposed to be extended to a government data center or similar registry of citizens
3. The technologies used on this project will have to conform to a FOSS license to ensure it is freely available to anyone

High Level Work Breakdown and Roadmap

The following table gives the high level breakdown and roadmap.

<i>Milestone</i>	<i>Quarter</i>			
	1	2	3	4
Disaster management domain research and requirements definition				
Sahana architecture (iteration 1: camp, people, org, request, security)				
Security and data privacy design and policy				
E-R model/database design (iteration 1: camp, people, org, request)				
Internationalization customization module				
Security and logging module				
Organization registry				
Camp registry				
People registry				
Reporting module (iteration 1: people, camp, org, security)				
Request management system				

Milestone	Quarter			
	1	2	3	4
QA and bug fixing buffer				
Sahana architecture (iteration 2: mobile, assistance, burial, damage)				
E-R model/database (iteration 2: mobile, assistance, burial, damage)				
Data consolidation module				
Mobile station development				
PDA and mobile phone access				
Assistance management system				
Burial module				
Damage database				
Reporting module (iteration 2: Assistance, burial, mobile, damage)				
Packaging and quick customization module				

Getting Requirements from NGOs and Gov

To ensure ROI and utilization of this system, the Sahana project will have to be presented and showcased to many different organizations to bring up awareness and visibility of the availability of such a project to handle disaster response efforts. Typical places where we hope to present this system are at disaster management conferences, FOSS conferences, NGOs (Red cross, CARE, Transparency international, etc), relevant UN groups and governments. NGOs and similar groups in particular can serve to provide a lot of domain-oriented requirements. It will also have to be hosted on its own website.

ASP and Voice portal extensions to Sahana system

Beyond the above-stated scope of the Sahana, there are two major extensions that could be considered to help better realize our objectives on quick deployment and easy accessibility. In terms of deployment, having the Sahana system under an ASP (application service provider) model on a website in which a new instance is deployed in an automated way to account for the latest humanitarian disaster allows for a very quick deployment. Secondly, having this centralized and secured on a known website under the authority of a well-known and certified humanitarian organization allows us to better secure and safeguard the sensitive data being entered in the system.

The second extension allows for better accessibility of the system in the field and for countries where there are low-literacy rates, which is to extend the access to the system through a voice portal.

Both these extensions require a significant amount of engineering and, if approved, the additional months with the defined team above is given below.

Task	Effort Months
Architecture and design for Sahana ASP version	1
Engineering automation of Website ASP version of Sahana	2
Setup of Sahana ASP Website	1
Voice portal interface module	3
Voice portal internationalization	2

Domain and Architecture Leadership Team and relevant roles:

1. Dr Sanjiva Weerawarna, Founder of Lanka Software Foundation, Author of the Web Services specification, OSI Board Member
2. Paul Currian, Information management consultant for humanitarian operations, UN WFP consultant, ICT4Peace, Author of HIC concept
3. Don Cameron, Regional officer NSW Fire Services, Emergency Management (NSW Police Academy)
4. Tom Worthing, Experience with disaster management in the defense background with a found on usable and efficient web interfaces
5. James Clark - Thailand Open Source representative, well known author of GROFF and EXPAT xml parser
6. Veriy - Indonesian Open Source representative
7. Chamindra de Silva – Sri Lanka FOSS contributor and project lead on Sahana phase I

References:

- [1] “The Open Source Initiative OSI”, (2005), (opensource.org), Available: <http://www.opensource.org/docs/definition.php>
- [2] “LAMP: The Open web platform”, (2001) (onlamp.com), Available: <http://www.onlamp.com/pub/a/onlamp/2001/01/25/lamp.html>
- [3] “Sahana Project”, (2005), (sourceforge.net) Available: <http://sahana.sourceforge.net/>
- [4] “Document A01 - About SourceForge.net?”, (2005), (sourceforge.net), Available http://sourceforge.net/docman/display_doc.php?docid=6025&group_id=1
- [5] “freshmean.net : About”, (2005), (freshmeat.net), Available: <http://freshmeat.net/about/>
- [6] “News - Opie - Open Palmtop Integrated Environment”, (2005), (opie.handhelds.org), Available <http://opie.handhelds.org/>
- [8] “Mozilla Developer Central”, (2005), (mozilla.org) <http://www.mozilla.org/developer/>
- [9] “CNO Main Page”, (2005), (cnosrilanka.org) Available: <http://www.cnosrilanka.org>
- [10] “Free/Libre Open Source Software”, (2005), (wikipedia.org) Available: http://en.wikipedia.org/wiki/Free/Libre_Open_Source_Software
- [11] The Cathedral & the Bazaar (paperback) By Eric S. Raymond ISBN: 0-596-00108-8