

RELIEF 10-01 After Action Report

Camp Roberts/NPS: 10-14 November 2009

STAR-TIDES

*Sustainable Technology Advanced Research –
Transformative Innovation for Development and Emergency Support
National Defense University*

Dr. Linton Wells II
Principal Investigator

John Crowley
Contractor to STAR-TIDES (acting as PM for Experimentation)

1 Dec 2009

Table of Contents

Background	4
Aerial Imagery	4
Field Data Collection	4
Integrating Systems	5
Preliminary Results	5
After Action Report: Project Insights and Lessons Learned	7
OPENAERIALMAP: Overhead Imagery Archive	7
TALKING PAPERS : Self-Describing Forms.....	8
UAV IMAGERY MOSAICS.....	10
<i>Lessons Learned:</i>	12
SAHANA DISASTER RESPONSE MANAGEMENT SYSTEM	12
SYNERGY STRIKE FORCE: Information Sharing and Virtualization	13
<i>Lessons Learned</i>	13
OVERALL LESSONS LEARNED	13
Technologies and Practices	13
Standards.....	14
Experimentation	15
Requests for February	16
Appendix 1: OpenAerialMap Future Directions	17
Attendees	17
Minutes.....	17
<i>OAM Reboot Technical Proposal: Catalog Quickstart</i>	17
Achieving Consensus, Initial Governance.....	18
<i>Communities of Interest (with many overlaps)</i>	18
<i>Exemplars for group management</i>	18
<i>Other organizations might also work as hosts for an OAM working group:</i>	18
<i>Project steering committee</i>	18
Licensing	19
Hosting	19
Data Distribution	19
Data Model Considerations	20
Action Items.....	20
Open Questions	20
Coda	21
Appendix 2: Talking Papers	22
Paper, Paper Everywhere	22
Yet Another “Last Mile” Problem	23
An Opportunity	23
Limitations of OCR.....	23
Self-Describing Paper.....	24
Talking Papers.....	24
Building Blocks	25
1. <i>Form Generator</i>	25

2. <i>Forms Reader</i>	26
3. <i>Data Scrubber</i>	27
How to Get Involved.....	27
Closing Thoughts.....	27
Appendix 3: Sahana Disaster Management System	28
Appendix 4: Synergy Strike Force	28

Background

The STAR-TIDES experiments at RELIEF 10-01 tackled three challenges to building a common operating picture:

1. Making effective use of aerial imagery.
2. Exploring new methods for field data collection.
3. Developing systems for integrating streams of information.

The STAR-TIDES team included representatives from Sahana; Open Mobile Consortium; InSTEDD; San Diego State University Visualization Lab; Open Solutions Group, Inc; Synergy Strike Force; Unthinkingly.com; and the National Defense University.

Aerial Imagery

In many parts of the world, building an overhead map of a disaster and getting it to first responders is hampered several factors:

1. **Cloud Cover.** In climates like Latin America, low cloud cover tends to have already obscured the ground before US satellites pass over the region.
2. **Slow Bureaucracies.** In all regions—including CONUS—imagery tends to flow through a slow bureaucratic process and rarely reaches responders within the first 24 hours.
3. **Access/License Restrictions.** Even when imagery reaches responders, it often carries restrictions on access which prevent sharing those images between mission partners (e.g., host country Red Cross/Red Crescent Society or NGOs based in other countries, like Doctors without Borders).

STAR-TIDES explored a rapid collection and distribution mechanism for an alternative source of imagery: small UAVs. The STAR-TIDES team experimented on the use of small UAVs to build a base map within hours of a sortie. It also assisted an open-source community of neo-geographers to build the technical architecture for a shared, open, and distributed repository of imagery to which responders could submit UAV (and toher) imagery and build a base map via collective action. This repository would be available to all mission partners and would make use (as much as is possible) of open licenses for imagery hosted therein.

Field Data Collection

While the use of laptops to collect data in the field has allowed for many advances in the type of information collected by staff who are assessing the aftermath of disasters, several challenges are hampering rapid analysis of the situation.

1. **Paralysis by Data Collection.** Because the vertical length of a web browser screen is nearly unlimited, organizations have devised forms with more than

100 fields per assessment. In Iraq (2003), one organization designed and deployed a form with 387 data elements, even after receiving pushback from the field. Such lengthy forms cause field staff to make decisions over what subset of those data elements actually get collected. These partial slices of data rarely create data sets that enable comprehensive and systematic analysis.

2. **Communications and Power Limitations.** Laptops work in environments where recharges are reliable and communications with a central operations center can be established. In the austere conditions after a major disaster, these assumptions cannot be made. The most reliable form of field data collection is still a clipboard, pad, and stubby pencil.
3. **Digital Divide.** If organizations are going to harness the local knowledge of host nationals, computer literacy cannot be assumed.

The STAR-TIDES team explored two options for confronting problems, albeit from different sides. The first, Talking Papers, extended the paradigm established by Walking Papers in the RELIEF 09-04 August experiments: using paper as advanced technology. Walking Papers prints maps on PDFs with QR codes that capture metadata about an area of interest (AOI), enabling field staff to scan annotated Walking Papers back into the base map for analysis. Talking Papers creates self-describing forms, where QR codes capture the data schema and OCR-ready field elements make optical character recognition (OCR) easier to perform in the field.

The second approach looked at low-power consumption Android phones to provide a means to encode forms using a touch-screen device that can communicate using low-bandwidth cellular connections.

Integrating Systems

Even when good data exists in the field, integrating those data into a common operating picture raises difficult problems. In many instances, the barrier between organizations—and especially between civilian and military organizations—prevents systematic flows of information between mission partners. When information gets shared, putting it into a format that can integrate streams of activities into a common operating picture has always raised difficulties—even with large military budgets, let alone on the shoestring ICT budgets of fielded NGOs.

The STAR-TIDES team explored the use of a distributed mesh of information sharing devices among mission partners to foster increased information flow. It also explored the use of the Sahana disaster management system to integrate those data flows into a composite picture of “who needs what, where,” including the integration of SMS messages to send and receive georeferenced requests for resources.

Preliminary Results

RELIEF 10-01 experimentation was another success. It led to the relaunch of OpenAerialMap (with 15 developers), the creation of an open-source effort to build Talking Papers (with 6 developers), the successful mosaic of UAV imagery using open source tools, the successful proof-of-concept for an Android-based field assessment

tool, and the successful deployment of Sahana to manage the experimentation at Camp Roberts TNT event, with over 200 people in the field.

These results can be applied in a number of geographic areas. These trials paid particular attention to:

- **Afghanistan**, because of its strategic importance
- **US Southern Command (USSOUTHCOM)** areas because of their support of experimentation and to evaluate capabilities that might be used in the PEAK (Pre-positioned Expeditionary Assistance Pit) JCTD, if approved, and
- **US territory** in light of support to RELIEF 10-1 from the **Department of Homeland Security (DHS)** and potential interest by the Administrator of the **Federal Emergency Management Agency (FEMA)** in restoring connectivity and information flow to citizens and local governments after disasters.

The remainder of this report details these successes, as well as the lessons learned during the preparation, execution, and analysis of the experiments at RELIEF 10-01.

After Action Report: Project Insights and Lessons Learned

OPENAERIALMAP: Overhead Imagery Archive

An imagery archive to exchange UAV and satellite imagery among HADR/SSTR mission partners

During the August RELIEF 10-01 experiments, the STAR-TIDES team processed a swath of imagery from eastern Afghanistan and prepared it for use by NGOs in the field. As a result of that work, the team identified a need for a shared, open archive of imagery which NGOs could access from the field and from which they could download a common set of satellite imagery. The team also hoped that this archive would be designed to be compatible with NGA's provisions for basic security and preservation of licensing restrictions.

Starting in late August 2009, the STAR-TIDES team mobilized an effort to reconstitute an archive which had once existed to fill a similar role: OpenAerialMap. Loosely defined, OpenAerialMap is the raster equivalent to OpenStreetMap: where OpenStreetMap provides vector tracings of roads and points of interest in an editable "Wikipedia of Maps," OpenAerialMap aims to provide an open archive of overhead imagery which can be used as basemaps and layers by other tools.

The STAR-TIDES team helped the community around the defunct project to identify the issues which had led to the collapse of the first iteration of their work and to secure resources which would put the project on a stable development basis. Importantly, STAR-TIDES left the governance issues and mobilization to the community itself, which—if it is to become a successful open-source software project—needs to own their project and have leadership over its direction. During the experiments, STAR-TIDES assisted Prof. Don Brutzman of the Naval Postgraduate School to host the conversations necessary for reconstituting the community.

Before, during, and after the RELIEF 10-01 experiments, the community itself was able to:

1. **Host an open conference call.** Fifteen (15) members of the OAM community participated in a conference call hosted from the Paso Robles Inn near Camp Roberts. Several members drove or flew to Paso Robles, CA to participate in the call, including Jeff Johnson, Travis Pinney, and David Bitner. The minutes from that meeting appear here:
http://openaerialmap.org/OpenAerialMap_Future_Directions.
2. **Create an initial governance structure.** The community decided to initiate their work with a project steering committee (PSC) elected by popular vote along the same lines as the OpenLayers project. Candidates would be nominated; if more than 5 candidates emerge, a vote would be held with two neutral referees from outside the project to count the votes. The five candidates with the most votes would serve on the PSC for a term of 1 year.

3. **Create a design for a scalable architecture.** The community decided to build a catalogue server infrastructure which will point to many independent storage nodes. This architecture enables redundant storage; possible advantages in speed by storing imagery closer to the requesting party (analogous to Akamai); and the opportunity for local storage nodes to enforce specific licensing or access restrictions (which will be critical for working with NGA NEXTVIEW imagery).
4. **Obtain donations to host and populate the bootstrap server.** NPS offered 10TB of backed-up storage, high-speed bandwidth, and access to supercomputer time for processing tiles.
5. **Start Recoding Catalogue Server.** Several developers started to modify the catalogue server code from the initial implementation to meet new requirements established during the conference call.
6. **Send project to RHoK in Palo Alto, CA.** Several developers participated in the Google/Microsoft/Yahoo/NASA Random Hacks of Kindness hackathon that was held concurrently with RELIEF 10-01, continuing work on the OAM catalogue server.

NB: The project team has also sustained momentum towards a vote to elect a PSC in late November. More information can be found on the OAM-Talk mailing list, which is accessible from the OAM Wiki at <http://openaerialmap.org/>. See Appendix 1 of this report for more information.

TALKING PAPERS : Self-Describing Forms

In a world that is making a slow transition from paper to pixels, the continuing need to use paper in the field creates a suite of problems. While paper is an excellent tool for data collection (it is cheap, light, works without recharging, offers a large and high-resolution canvas, and even is legible when wet or torn), it is a poor tool for transporting data from one place to another (in bulk, a mass of paper forms is heavy, flammable, and very time-consuming to transcribe). Preprinted paper forms are also inflexible: they cannot ask new questions that emerge as a result of successive surveys. As a result, paper introduces two time delays: a delay between data collection and data analysis, as well as a delay between analysis and new surveys. The STAR-TIDES team examined a novel means of reducing both delays in today's field data collection regimes.

Borrowing from the technology that created Walking Papers in the August RELIEF 09-04 experiments, the STAR-TIDES team (lead by Robert Kirkpatrick, the chair of the Open Mobile Consortium) built a OCR-readable paper form that can be generated on the fly and which embeds its data schema into a set of QR codes that run along the top and bottom of the form in a strip:

**Missing Person Entry
(Reporting Person)**

Identity

If you have any information about missing person's identity numbers, Please include them since it will be helpful for finding the person. Otherwise please ignore.

Identity Card Number [REDACTED]

Passport Number [REDACTED]

Driving License [REDACTED]

Basic Details

Please enter any of the following details on the person:

Full Name [REDACTED]

Family Name [REDACTED]

Local Name [REDACTED]

Date of Birth [REDACTED] Age Group [REDACTED]

Gender [REDACTED] Marital Status [REDACTED]

Religion [REDACTED] Race [REDACTED]

1 of 3

UPLOAD IMAGE AT [HTTP://WWW.TALKING-PAPERS.ORG](http://www.talking-papers.org) OR EMAIL TO FORMS@TALKING-PAPERS.ORG

This technology, dubbed Talking Papers, will enable field workers to 1) generate forms that match the current data collection priorities, 2) collect data on paper, 3) scan the completed forms back into a database via OCR, and 4) modify the form elements to ask questions as they arise. As a result, paper can become a data collection tool well-suited to the austere conditions while not delaying the analysis of emerging trends.

The team was able to harness the energy of at least five developers to build the tool, including the inventors of Walking Papers, Mike Migurski at Stamen Design and Josh

Livni at Umbrella Consulting, as well as interaction designer Chris Blow at Unthinkingly.com.

For more information, see: <http://humanitarianetech.com/2009/11/16/talking-papers-a-world-without-data-entry/>. The code repository is available at <http://github.com/unthinkingly/talkingpapers> and the mailing list is available at <http://groups.google.com/group/talkingpapers/topics>. See Appendix 2 of this report for more information.

UAV IMAGERY MOSAICS

For fielded organizations, quickly obtaining overhead imagery of a disaster area is still a largely-unfilled dream. Satellite imagery often takes days or weeks to permeate down to the level of field staff working in forward bases; even then, the imagery may be put under access restrictions which make it impossible to share with local partners, including members of settlement- and village-level government officials in the host nation.

UAVs offer great promise to fill this gap, providing teams inserted into local areas with a tool to map an area sortie by sortie, using flight patterns that optimize coverage and provide ultra-high resolution imagery of points of interest. However, in the civilian space, this technology is rarely used; even in cases where it is available, UAVs are still treated as a proof-of-concept device. Few organizations have a concept of operations or the requisite software to weave together a series of still photographs into mosaic that can be used as a basemap. To the knowledge of the STAR-TIDES team, no organization has explored the use of multiple simultaneous small UAVs to quickly map a region as a collective.

At RELIEF 10-01, the STAR-TIDES team (lead by two neographers from the Open Solutions Group, Inc, Jeff Johnson and Travis Pinney) explored the most basic functionality for field use: the use of UAVs to create a base map of a small area. Working with data collected by Prof. Kevin Jones of NPS, two geographers created a set of Python scripts to map UAV still images from a gimbaled COTS camera on the RASCAL platform onto a 1m resolution base map of Camp Roberts. This technique took metadata embedded into each image and used it to calculate the placement of the approximate image on a map. These metadata were as follows:

Event details follow:

```
Rascal UAV ID: 3
  Date           : 00/00/2009
  Time           : 00:00:00.00
  UAV Latitude   : 35.736678 degrees
  UAV Longitude  : -120.798284 degrees
  AGL Altitude   : 417 m
  MSL Altitude   : 687 m
  HAE correction : -7 m
  Heading        : -62.7 degrees
  Speed          : 28.9 m/s
  Roll           : -2.0 degrees
  Pitch          : -5.0 degrees
Target:
  Target Latitude : 35.736670 degrees
```

Target Longitude : -120.798297 degrees
Target Elevation : 270 m
Camera:
Model : Canon PowerShot G9
Lens : 7.4 - 44.4mm (35mm equivalent: 33.8 - 202.9mm)
Exposure:
ISO : 200
Aperture : 4.5
Shutter Speed : 1/1250
Settings:
Resolution : 4000 x 3000
Quality : Superfine
Image Stabilization : Off
Focal Length : 7.4mm (code:)
Focus Distance : 1.92 m
Hyper Focal Distance : 1.85 m
Depth of Field : 0.92 m to infinity
Photo:
Center Latitude : 35.736389 degrees
Center Longitude : -120.797597 degrees
FOV (degrees) : 54.2 x 42.1 degrees
Projected Pixel Size : 10.8 cm
Calculated corner coordinates:
Lon1 : -120.796283 degrees
Lat1 : 35.734212 degrees
Lon2 : -120.800544 degrees
Lat2 : 35.736000 degrees
Lon3 : -120.798910 degrees
Lat3 : 35.738566 degrees
Lon4 : -120.794649 degrees
Lat4 : 35.736778 degrees

Based on the bounding box of the Calculated Corner Coordinates and the heading of the aircraft the time of the photo, the team created a tool to mosaic imagery within several minutes of image receipt.

The resulting product appeared with many tiles up to 15 degrees off alignment (see image below, note airfield in lower left in two very differently rotated tiles):



Upon investigation with Kevin Jones, the team discovered that the Inertial Measurement Unit (IMU) of the UAV in question (RASCAL) based its heading on a GPS compass instead of a calculation that incorporates the potential crabbing of the aircraft off the GPS heading (as would occur in a crosswind). This discrepancy accounts for the inaccuracies of the script. It could likely be addressed through improvements to the IMU of the UAV as well as through complex imagery post processing. The STAR-TIDES team will investigate further fixes in later experiments.

Lessons Learned:

1. **A GPS Compass is not good enough:** The metadata embedded into a UAV still image must account for the yaw, publishing how many degrees off center the airframe of the bird may facing relative to its direction of travel.
2. **Pre-stage imagery to the field.** The STAR-TIDES team experienced periods when our Internet connection failed. During these periods, WMS servers, Amazon EC2, and other imagery were not available for use by the team. STAR-TIDES recommends that any fielded organization should provide tools that ensure that field staff do not have to download imagery over thin pipes. Instead, organizations should load their machines with one or more good basemaps as well as other imagery. They should also make sure that imagery is in a fully open format that can be shared freely among partners and opened in any application (that is, make sure the imagery is not locked to any application or vendor, like Falconview, Google, or ESRI).

SAHANA DISASTER RESPONSE MANAGEMENT SYSTEM

The Sahana team partnered with members of the Trinity College Humanitarian Free and Open Source Software (HFOSS) program for two weeks of experimentation, supporting both RELIEF 10-01 and TNT 10-01. Because their experiments were directly

sponsored by NPS, they have submitted their own After Action Report, which is included as Appendix 3.

SYNERGY STRIKE FORCE: Information Sharing and Virtualization

With growing complexity of software and even more complex configurations between multiple pieces of software, deployment by individual installations on specific hardware platforms and operating systems can be both costly and slow. The STAR-TIDES team (lead by Todd Huffman of Synergy Strike Force, or SSF) examined how to use virtual machines to reduce the time for deploying complex software quickly.

Synergy Strike Force is a private volunteer organization that is assembling a suite of minimal essential software for HADR/SSTR/COIN operations. This tool set is intended to foment information sharing across the civilian-military divide, building the capacity for unity of effort when there is not unity of command.

Building on work performed during RELIEF 09-04 in August, the SSF team explored the use of virtual machines to quickly deploy multiple packages of open source software. The suite of software examined during RELIEF 10-01 included a document management system, a conferencing system (with VOIP, chat, and desktop sharing), and a wiki. See Appendix 4 for more information, including an ORD and TRD that evolved from the work that SSF performed at CP Roberts.

Lessons Learned

1. **Support multiple virtual-machine client-software applications.** While the current tools are built on JukeBox and VirtualBox, the USMC has an enterprise license for VMWare. It would be beneficial to provide the information sharing system package of software on virtual machines that run across all these platforms.
2. **Private subnets with local domain name servers simplify deployment.** Expecting humans to remember numerical IP addresses is not prudent. When disconnected from the Internet, it would be useful to name individual machines with human-readable domain names for configurations that require communications within and across organizations in the field. This problem will become more acute as 8220.02 (and network sharing between civilian and military organizations) becomes more prevalent.

OVERALL LESSONS LEARNED

Through STAR-TIDES work at RELIEF 10-01 and conversations with NPS, USMC 1MEF, and other organizations at the experiments, the following lessons-learned emerged:

Technologies and Practices

1. **FEMA and SOCOM both Believe that “Communications First” Extends to the Affected Population.** DHS/FEMA and SOCOM have a shared problem: how to raise the level of communications infrastructure not just for the responders, but for everybody—with the expectation that the affected population owns many resources

by which they can bootstrap themselves from disasters. FEMA Administrator Craig Fugate wants to enable affected population after a disaster; so does SOCOM with Indigenous Support Kit (ISK). That said, while DHS/FEMA and SOCOM have important knowledge to exchange, no formal structure exists to facilitate interaction on this shared problem. This interaction should be fostered on an immediate basis.

EXAMPLE: After previous deployments to Iraq, 1MEB at CP Roberts had amassed deep knowledge of application of UAVs and open source software to managing complex emergencies. FEMA had knowledge of tools and practices for managing civilians and using civilian technology of which 1MEB was unaware and could have used during its campaign in Iraq. It took a hotwash at RELIEF 10-01 for STAR-TIDES to surface these shared, cross-agency interests.

2. Reducing the Risk of Adopting New Tools Requires Careful Thought. Many First Responders are not aware of what tools and technologies currently exist and how to evaluate them. As a result, they will tend to “do without” rather than invest in an unknown tool. This risk-averse approach is particularly present in organizations with low discretionary funding (1-2%). If DHS and DoD wish to give confidence to these responders to make intelligent acquisitions, they will need to lower the risk of adoption of new tools and methods. STAR-TIDES will need to add data relevant to tool adoption if it expects responders to make informed choices that involved perceived risks.

Standards

1. **Data-level Interoperability: Common Data Standards are Critical.** Data-level interoperability provides the basis for effective information sharing. It is critical to establish common data formats between all the entities that are coming together for HADR, SSTR, and COIN operations, as well as all the organizations that come together for RELIEF and TNT experiments. It acknowledges that each organization will continue to use its own suite of applications, but provides for a *lingua franca* between those applications at the level of data feeds. Where possible, organizations should publish their data schemas so that other partners can validate information flows against those schemas and write translation adapter for those data flows. Organizations should work collectively to name things similarly (ontology) and to classify things in similar ways (taxonomies).

For experimentation, it will be critical to develop a common dataset for all the participants to use, share, extend, and test against. STAR-TIDES may also need to build bottom-up support for standards around experimentation at CP Roberts. Such support is already emerging around OAM, OSM, SMS-Field Reports, and Talking Papers.

2. **Publish Standards for First Responders.** DHS needs to know which standards to push in ORDs for Situational Awareness Software packages. The experts on

the STAR-TIDES team recommended the following standards be incorporated into future operational requirements documents:

- WMS
- KML
- GeoRSS
- Tile Mapping Service. TMS (but also look into WMS-C)
- WFS
- CAP
- XMPP
- OpenID
- OpenLDAP

The team also recommended that DHS and DoD follow the World Wide Web Consortium (W3C) Emergency Interoperability Group for other standards that will emerge around HADR and SSTR.

Experimentation

1. **New Invitation Process for Participants.** For future RELIEF experiments, NPS and NDU will process invitations as part of the biweekly conference calls. Either NPS or NDU can suggest an avenue of research and suggest some companies who might be good to invite. The other partner will vet those companies and ensure that the partnership and invitation process is proceeding according to established rules, especially for commercial entities and for intellectual property created. NDU will also investigate creating an open BAA for Camp Roberts.
2. **Stable Funding is Required for Expanded Experimentation.** The RELIEF experiments at CP Roberts will remain an ad-hoc effort until a sponsor is willing to commit multi-event funding and organizers can focus on building community support amongst first responders and international NGOs instead of dealing with funding/administration issues.

In addition, funding must be made available to non-profit entities that wish to participate and can contribute their tools and people to solving core problems in RELIEF's domain. Non-profits, NGOs, and universities lack funds to send staff with cutting edge tools to CP Roberts. Even when they are able to send staff, they have funds only for one trip. If STAR-TIDES is going to participate in these experiments, and if it expects to create a learning dynamic, where the lessons from previous experiments are incorporated into the next round of experiments, there needs to be a fund to enable organizations to send field staff, engineers, and researchers to CP Roberts.

3. **Branding can improve.** For Twitter, it is important that we start using the brand RELIEF. Staff will requests that participants use the Twitter hashtag of #RELIEF1001 or something similar, as #camproberts is neither our only location nor the name of the event.

Requests for February

The team requested the following resources for February:

1. Reliable high-speed access to the public Internet.
2. A small "continuity" server to meet the following needs:
 - Store experimentation tools and data between events
 - Host a DNS Server for human readable server aliases
 - Host a comms server for collaboration with minds external to site and btw sites
 - Host an onsite Wiki where individuals can collaboration before, during, and after the events from anywhere on the planet. This wiki would need to be open the public, and may need a solution for local (CP Roberts) use during periods of disconnected networking. Also, OpenAtrium is not working well for the team; RELIEF may need to consider other collaboration platforms.

All of these servers could be virtual machines hosted on 1-2 physical servers. Servers could likely be very small, inexpensive boxes, like a MacMini or small desktop from Staples.

Appendix 1: OpenAerialMap Future Directions

On 12 November 2009, nine of us met face-to-face in Paso Robles California USA, and six more folks dialed in. Discussion lasted about 2 hours. This was an excellent meeting with lots of great ideas and contributions.

Attendees

- David Bitner, OSGeo/Sahana
- Jessica Block, SDSU
- Don Brutzman, NPS
- Jon Crowley, STAR-TIDE
- Stefan de Konink, OpenStreetPhoto
- Schuyler Erle
- Cristiano Giovando, EC-JRC
- Todd Huffman, STAR-TIDE Strike Force
- Jeff Johnson, Open Solutions Group
- Don McGregor, NPS
- Hugo Meiland, Leiden University
- Jill Olen, San Diego
- Marc Pfister, ENPLAN
- Travis Pinney, Open Solutions Group
- Charles Schmidt, Terrapan Labs

Minutes

OAM Reboot Technical Proposal: Catalog Quickstart

Schuyler first discussed his OAM Reboot Technical Proposal and the existing SVN repository at <http://svn.openaerialmap.hypercube.telascience.org>

Catalog server does seem like most important initial task facing us.

- We are building a Web application
- This will be front-end for submitting and retrieving OAM data
- Catalog capabilities will be cross-connected to other GIS server capabilities
- Hosting site must offer reliability, security, up-time, capacity and access
- Can we compare software alternatives to accomplish this?

Application Frameworks: two candidates seem pertinent.

- Django, written in Python, familiar to GIS community, currently GeoDjango plugin is superior solution
- Rails, written in Ruby
- Zend, written in PHP (probably not ideal since it is lacking GEO specific code (Hugo))

Functionality seems approximately equivalent. Each can talk to an SQL relational database on the back end. Might need a REST interface.

Might need to swap in a different application layer later, so it might be smart to quickly put up something workable as quickly as possible to learn lessons and reach conclusions, throwing away first-round application isn't necessarily bad.

Random Hacks of Kindness is a big opportunity. It would be cool to have a prototype catalog server up and running by next Monday for discussion purposes. Perhaps other versions will be built later using other application frameworks later for comparison - nothing wrong with that. Communications will be via **Error! Hyperlink reference not valid.**

Achieving Consensus, Initial Governance

Public discussion is REALLY IMPORTANT. Back-channel discussion is helpful for initially figuring out some things, but more critical is that the knowledge gets disseminated and considered by the group.

Achieving group consensus through open discussion is the most critical thing we can do. Posing questions, options, tradeoff alternatives, etc. means that people can understand what is going on. If we act with mutual trust and openness during these critical reboot efforts, we are setting the pattern for effective progress.

Figuring out who is part of the OAM developers and looking at other exemplars will help us figure out how to best establish longer-term patterns for effective cooperation and self governance.

Communities of Interest (with many overlaps)

- Geographic Information System (GIS) developers and users
- Remote Sensing and scientific users
- Emergency Management, support for first responders and disaster relief
- Web developers and users

Exemplars for group management

- Open Street Map Foundation
- OpenLayers steering committee

Other organizations might also work as hosts for an OAM working group:

- Open Source Geospatial Foundation (OSGeo)
- Open Geospatial Consortium (OGC)
- Web3D Consortium

Project steering committee

- Keep coherent shape to many concurrent efforts
- Role includes ensuring that consensus is followed when possible, variations are certainly OK but should be documented and further discussed

- How to bootstrap? Mailing list self-nominations, initial volunteers for interim period, voting?
- Both coders and community domain expertise are needed for effective decision making
- Committee primarily guided by consensus, lack of consensus usually means more work is needed (and forced decisions often turn out to be ineffective anyway)

Licensing

Our goal for OAM is to accept and serve open-source code and open-access data. Several important questions pertain.

- What open-source licenses are acceptable for code?
- What open-access licenses are acceptable for data?
- Public Domain
- Creative Commons
- Likely different licenses for code and for data
- Attribution and redistribution requirements
- Can we prevent malicious misuse or repurposing of data?
- How do we verify identities of submitters?
- Can we embed or link license metadata within image files themselves (e.g. JPEG2000 or GeoTiff)

Enough is known about our licensing goals at this point that we can proceed with the prototype. Further discussion and documentation will be needed.

Hosting

- Telascience server SDSU-UCSD - John Graham
- NPS - Don Brutzman and Don McGregor
 - NPS is both .EDU and .MIL. The .EDU side is quite open, while .mil can be restricted.
 - Open aerial map is international in nature. Our priority is open, international access.
 - NPS has connections to CalREN, has ROCKS cluster. The cluster is not publicly exposed, but can act as a behind the scenes resource
- Other sites can also host the data

Data Distribution

Goal is to create a distributed data model with multiple sites providing various degrees of caching. This is important discussion topic for this weekend.

- Updates might be via hard-disk (sneakernet) but direct connection over network needs to be possible.
- Also need to be able to download/cache/disconnect data on demand, in either modest or large quantities.

Data Model Considerations

There are many relevant data standards available. Nevertheless there is also a need to keep the number of `_required_` standards to a bare minimum, so that minimum functionality can be maintained.

- Critical standards bodies and groups for geospatial metadata include
 - Federal Geospatial Data Committee (FGDC)
 - International Organization of Standardization (ISO)
 - ISO 19115-2:2009
 - ISO/TS 19129:2009
 - INSPIRE in Europe
 - INSPIRE Metadata Guidelines
 - Open Geospatial Consortium (OGC)
 - Earth Imagery
 - Metadata
- What are critical data standards that we can start with?
 - Support for GDAL: gdalinfo
 - Dublin Core for document metadata
 - Geography Markup Language (GML)
 - License name, link
 - Need to list others here, many exist
 - Are there any that would help initial catalog-server standup effort?
 - What metadata conventions does OpenStreetMap follow? (Reportedly no metadata requirements are used, need to confirm...)
- Additional data standards are OK but probably considered extra
- What are criteria for including a standardized set of metadata on the short list of required inputs?
- Might even want to allow well-formed but unstructured user-defined metadata?

Action Items

Likely there are plenty more action items that will flow from all of the things we discussed. Here are some.

- **Error! Hyperlink reference not valid.**
- Integrating these points into the OAM Reboot Technical Proposal
- Schuyler will post proposal about how to establish an initial interim Project Steering Committee

Open Questions

- What else do we need to consider?
- How will the massive amount of data be stored. The previous infrastructure couldn't handle all the available data now. How will OAMv2 address this hosting and storage (distribution) wise? --Skinkie 17:25, 12 November 2009 (UTC)

- License-based serving; if a user is only interested in CC-BY, it should only serve all compatible licenses under it and not CC-BY-SA (above it). --Skinkie 17:30, 12 November 2009 (UTC)
- Once an initial catalog server has been stood up, what do we do next? Another workshop meeting perhaps?
- Might a monthly teleconference/IRC meeting be helpful? Maybe in January...

Coda

This was a productive session today. A lot of productive discussion occurred. We hope this weekend's code-a-thon goes well!

Additional edits to the minutes are welcome to ensure that they are accurate and complete.

New points and ongoing discussion on these topics can best occur on the OAM talk mailing list.

Thanks to all participants, and thanks for considering these powerful possibilities.

Appendix 2: Talking Papers

Humanitarian Data Collection 2.0

Robert Kirkpatrick, Chair of the Open Mobile Consortium and lead of Talking Papers, published the following blog post on 16 November 2009 at <http://humanitarianatech.com/2009/11/16/talking-papers-a-world-without-data-entry/>.

Last week at [Camp Roberts](#), entrepreneur [Todd Huffman](#) was kind enough to take me on a tour of [Walking Papers](#), a remarkable service that allows users to print out paper maps, annotate them manually, upload them into [OpenStreetMap](#), and use the annotations to transcribe new content. It's like digital tracing paper. Walking Papers is a brilliant idea in its recognition that paper – like it or not — still has an important role to play in field environments.

What really caught my attention was that the paper forms Walking Papers emits encode map quadrant coordinates, as well as a unique identifier, in a [2-D barcode](#) that is used to process annotated maps once they've been scanned and uploaded. When a map is uploaded, Walking Papers is able to read the barcode and plot the location on the globe to which the scan corresponds. Although it's not yet possible for Walking Papers to decipher my annotations automatically, the barcode is at least machine-readable: once the scan has been uploaded, I can take it from there to transcribe what I have drawn. This imaginative and insightful approach got me thinking about a related problem I've been keen to address for some time: data entry. How can we use paper as a more effective channel for information flow during and after humanitarian emergencies?

[Paper, Paper Everywhere](#)

In every disaster zone and every rural development environment where I've worked, paper is still king when it comes to collection of structured data, from population needs assessments, to tracking inventory stocks levels, conducting health surveys, filing situation reports, logging security incidents, and in general maintaining shared awareness of the situation unfolding on the ground. In spite of more than a decade of work by literally hundreds of organizations developing PDA-based data collection systems, the default option in the deep field remains unchanged: print out a form, take it to the field, fill it out with clipboard and stubby pencil, bring it back, and enter the data manually at a computer.

One day, hopefully soon, we won't need paper in the field. But that day is still years away. There are many reasons for paper's continuing status as the tool of choice for field data collection. It's cheap. It's light. It's compact. It doesn't need recharging. It doesn't need Internet access. It's a familiar – and spacious — form factor. It works in hot weather or cold. It can be read under bright sun. It's not affected by dust. Yes, it fares poorly if it gets wet, or torn, or smudged, and options for sensor integration and data validation are...lacking, but on balance, as a data collection tool beyond the edge

of the network, it still has a lot going for it. As a tool for data **transport**, however, particularly in crises where time is of the essence, paper is inefficient and ineffective.

Yet Another “Last Mile” Problem

If you depend on the data collected on paper forms to understand the needs of vulnerable populations and make decisions that affect their welfare, **paper is the weakest link in your information supply chain**. At virtually every stage in a paper-based process, there is room for human error to alter or lose critical data: when it’s written down, during transport, when it’s read, when it’s entered into a database. Paper is a fragile medium to begin with, but paper in the hands of hot, tired, busy, stressed-out relief workers in the chaos of a major disaster is fraught with problems. As long as paper is used for data collection, error and data loss will continue to reduce the effectiveness of humanitarian coordination, and unless someone invents self-validating paper, it’s hard to see ways that technology can help here anytime soon.

An Opportunity

There is, however, one shortcoming of paper that we might be able to address today. Virtually everywhere in the world of relief and development, completed paper forms accumulate in piles until someone has the time to enter the data manually into a spreadsheet, database, or other application. Data entry is not only a juncture where errors tend to be introduced; it’s also the point that tends to contribute most heavily to latency in the flow of humanitarian information. When critical information needed to match needs to resources reaches decision-makers too late, coordination breaks down, further delays are introduced, resources are misallocated, and too little arrives too late to help a population in need. Components of a potential solution to this data entry problem already exist, though no one seems to have solved it decisively. Before I suggest where we might go, I need to explain why current tools haven’t filled the gap.

Limitations of OCR

Optical Character Recognition (OCR) technology, for example, has been around for decades and has improved markedly in recent years. [Sahana](#), who brought a team to Camp Roberts, have already done some [excellent work](#) in configuring their disaster management system to emit OCR-friendly forms, and I am convinced that such approaches have tremendous potential to increase the viability of OCR in the field and the quality of the data captured. But there’s another reason OCR along won’t eliminate the need for data entry in humanitarian work the way it has, say, for many of the forms we complete in a non-crisis setting. A major limitation in applying OCR to paper forms in a humanitarian context is that the underlying **schema** of the data being collected is itself in a state of constant flux.

Emergencies are by their very nature dynamical systems characterized by emergent effects. Weather, disease, and natural hazards may worsen conditions without warning. Poorly understood needs – or poor communications – may have secondary effects that change the situation on the ground dramatically. Populations affected by the situation may respond in unforeseen ways – constructive and destructive – in ways

that alter both availability of resources and their need for them. Political decisions, news reports, and the choice of a single word may all change the course of events. What one thought one needed to know yesterday may no longer be important to ask today, or it may have been the wrong question all along, and as the response moves from critical intervention to mitigation and recovery, needs keep changing.

As a result of this dynamic, the forms designed to assess population needs at the outset of a response soon become inadequate. Questions must be added. Others must be removed. The schema of the data being collected has changed, impacting form and database design. A few days later it happens again. And again. And layouts change, as does the wording of questions. In many cases, updated or entirely new forms are designed, printed, distributed and collected in the field. Even if OCR could be used to extract data from these forms with 100% accuracy, it would do little for a decision-maker looking to make sense of the data, because this data is organized according to an unfamiliar schema that emerged at the edge of the network.

Self-Describing Paper

Walking Papers, however, suggests a way forward. That little barcode in the corner in effect contains a machine-readable schema for the map annotations, and it got me thinking about an [article](#) I read several years ago which noted that PDF-based forms could potentially encode their schemas automatically within 2-D barcodes. I find this idea fascinating. Print such a PDF, and you have a paper analog of XML: a self-describing document, machine- and human-readable, that contains both data and the schema describing that data. It's a paper form that tells you what it is. After a chat with [Mike Migurski](#) of Walking Papers, I'm code-naming the concept "Talking Papers", and I'm hoping to get a team together to work on making it a reality.

Talking Papers

Imagine the following scenario. There has been a major earthquake, and you're a nutrition expert working in the Food Security cluster in a makeshift office near the center of the affected area. You design a Household Nutrition Survey form on your laptop, pair it with a Walking Papers map of a village, print out 100 copies, and hand them out to a few trusted local volunteers to take house to house. As completed forms come back to you, you quickly scan them into your PC – no waiting for time to perform cumbersome manual data entry. Auto-magically, the data and metadata are extracted from the form and – Internet access permitting — uploaded directly into an online, collaborative environment where you and your colleagues review, correct, and validate the data against its schema. Once scrubbed, the data moves on to the next step in the supply chain: some download it in one or more standard formats, while others publish it into online repositories such as [Freebase](#), [GeoCommons](#), [DevInfo](#), [Sahana](#), [Mesh4X](#), [RapidSMS](#), [GATHER](#), etc. for mapping, analysis, and sharing.

Building Blocks

I've bounced this idea off of a few folks already, including [John Crowley](#), [Matt Berg](#), [Chamindra De Silva](#), [Todd Huffman](#), [Ed Jezierski](#), and [Chris Blow](#). We've agreed that making this work will require, at a minimum:

- 1) a tool to create printable forms,
- 2) a tool to read uploaded scans of completed forms, and
- 3) a tool to review, scrub, and publish data once it has been extracted.

Ideally, building, reading, and scrubbing features should be available offline, since Internet access is a scarce commodity in places where Talking Papers would be most useful, but it probably makes sense to get an online, browser-based version up and running first to get user input as quickly as possible. I think each tool should exist as a completely separate service, as there may be other uses for such capabilities. Where existing tools can be modified to address the requirements described above, I'm all for it.

Below are a few initial thoughts on building blocks.

1. Form Generator

This should be a user friendly online tool with a drag-and-drop interface that allows users to design text-entry-friendly, OCR-friendly forms with an option to export to PDFs. The tool would encode a serialized version of the schema in a supported standard format (e.g., [Turtle](#), [XForms](#)) within a band of high-capacity 2-D barcodes directly on the form. The barcodes should be duplicated across both the top and the bottom of the printable form for redundancy. Fields on the form, in addition to human-readable text labels, might have tiny machine-readable labels – perhaps also in barcode format – that associate the values that follow with data elements in the encoded schema. When a form is created, the designer would specify a default URL for the data-scrubbing workspace to which scans will be uploaded and processed, so that URL could also be encoded in the barcode – making Talking Papers not only self-describing, but self-routing. Ideally the tool would be able import schemas in standard formats generated by other tools and let users work with those as a starting point for form layout.

Here is a mock-up of a Talking Papers form I've annotated in red, based on a Sahana OCR form:

Scrubber. This payload should also include the UID of the associated Walking Papers map, if any.

3. Data Scrubber

This online application would help users clean up data sets within collaborative workspaces where they can review, edit, and publish data processed by the Forms Reader. A simple data grid UI would be a great start, and Google Spreadsheets would probably take us part of the way there. Data successfully extracted should be displayed in rows, with columns corresponding to each field in the schema. Some visual indication – perhaps coloring or shading? — could indicate where content was suspect, or could draw the eye to blank areas where OCR failed entirely. It might also be helpful to have a feature that detects and highlights or clusters duplicate entries. Each row should contain a link to the original scan to assist users in inferring the original intent of the individual who completed the form (e.g., to review content that OCR could not interpret, as well as marginalia and other annotations). Once the user is happy with the content, hooks should be provided to allow him or her to download the data in common formats or push the data into a variety of repositories.

How to Get Involved

Chris Blow has generously agreed to contribute to the design process; he has already stood up a repository at <http://wiki.github.com/unthinkingly/talkingpapers> and has begun working on scenarios. We need your ideas, suggestions and concerns. We need designers, developers, testers, and user-practitioners willing to test this system in the field and help us shape it into something genuinely useful. Substantial early user input, an agile, open-source, collaborative process, support for open data standards, and well-designed mashup-friendly APIs will be critical.

If you are interested in contributing to this effort, please contact me.

Closing Thoughts

Most of the concepts underlying Talking Papers are not new. Many of the required building blocks already exist in some form. But these capabilities as far as I know have never been brought together in a simple, flexible implementation that will actually work in the humanitarian field. What if we could design a system that generates and reads such forms, creating a seamless bidirectional bridge between paper and applications that takes you from data collection to data scrubbing in one hop – skipping data entry entirely? Can you imagine a paper-based XForms client? I believe strongly that this kind of technology could help to streamline the information supply chain in humanitarian operations dramatically, allowing those who depend on such information to save lives.

Appendix 3: Sahana Disaster Management System

Sahana Software Foundation

RELIEF Quick Look Report

November 25, 2009

Title of Project: Sahana Free & Open Source Disaster Management System

Principal Investigator: Mark Prutsalis, President & CEO, Sahana Software Foundation

Qualitative Results:

The following is a list of the qualitative results we achieved at RELIEF 10-1 (and TNT 10-1). Without exception, Sahana did not have any of these capabilities when we arrived on site on November 10, 2009, and were developed by our onsite team supported by a virtual team of developers who contributed code and ideas from remote locations:

- We used the Organization Registry and Volunteer Management System to track all of the persons and organizations and experiments taking place at RELIEF 10-1. This was a valuable demonstration of capabilities.
- We developed integration with webcams for taking digital photos of registering individuals/volunteers.
- We used a \$400 (two years ago price) eeePC/Netbook as our primary data collection server as a proof of concept of the low barriers to entry of Sahana.
- We configured Sahana OpenLayers to pull in WMS layers from mosaic-ed and geo-referenced UAV and satellite imagery of the Camp Roberts environs.
- We had a lot of discussions about using open standards for data exchange with government and non-governmental participants. We see Sahana as being a leader in this area.
- We had the opportunity to brief FEMA Administrator Craig Fugate on Sahana and the benefits of open source / open standards in the humanitarian/disaster response application field during a conference call with technology companies from Silicon Valley
- We configured an SMS gateway running on a Windows-based server using Cygwin and SMSTools and a Nokia 3220 phone to send and receive SMS messages from Sahana.
- We developed an Android application to send in structured SMS messages to Sahana with embedded GPS coordinates. We also developed the format for a structured SMS message to be sent from any cell phone to Sahana.
- We developed a system to register a user name to a cellphone number.
- These inbound messages (from Android app and simple cellphone) are now processed by Sahana to plot points on the situation map using the DHS

symbolset of incident information based on feature class reported by the Android and SMS message.

- We utilized a lot of DHS and NIMS/ICS terminology throughout the experiments to demonstrate to a US-audience how easily Sahana can be configured to be “NIMS-compliant”. This will eliminate an important barrier to entry for US government entities at all levels of government. As part of this, we have blueprinted the requirements to have Sahana automatically generate one of the ICS forms that first responders are required to fill out
- We developed the ability to poll the Sahana server and pull information about the last known location of a registered Sahana user, the last report(s) sent in by any Sahana user, or a keyword search of all points of information plotted on the situation map. In the future, we’ll want to develop more powerful means to pull data onto handheld devices.... both through SMS and where the Sahana server can be accessed directly by an app.
- We created a KML feed from Sahana data such that we could call it up in Google Earth. This is very attractive to a lot of potential users who use google earth to aggregate data from different feeds.
- We participated in an integrated field experiment with observers sending in reports via SMS and Android back to a command center that was utilizing all of these capabilities to get situational awareness of an event. Sahana, of course, outperformed a lot of other systems that had had a lot more financial resources put into them.

Any Initial Quantitative Results:

Our results are mainly qualitative in nature; in the future, we plan for more quantitative (load-testing) of Sahana (and the SMS system) in future RELIEF events.

Lessons Learned:

- Open Source / Open Standards for data exchange provides incredible flexibility for data integration and sharing. By maintaining our focus in this area, we are able to consume data from most other producers of visual/spatial and textual data as well as publish Sahana data so others can utilize it in their own systems. We support the widespread adoption of WMS, WFS, KML, Geo-RSS, EDXL and CAP.
- The SMS capabilities gives Sahana the ability to serve as an incredibly powerful crowdsourcing and assessment application (akin to Ushahidi, Geochat, DevelopmentSeed and others), and combined with Sahana’s native disaster management capabilities, gives Sahana users powerful situational awareness not available in other systems.

The following are capabilities that we have identified for priority development at future RELIEF and TNT events:

- We need to be able to package pre-configured deployment-ready instances of Sahana on virtual machine images, which will make Sahana available instantly after a disaster for use by response teams and the public.
- We need to run throughput tests on the SMS system and gateway; we experienced some lags (20 minutes or more at times) in SMS sent from one provider (Verizon) to be delivered to the provider of our gateway (T-Mobile). These should be documented in a more formal experiment so we can document appropriate recommendations for packaging Sahana for deployments.
- We need to expand our data collection capabilities to include enhancements to the Android app and other smartphones (the iPhone, Blackberry and other platforms / mobile operating systems), and from simple web-enabled cellphones using java-based forms (X-forms / javarosa). This would include using drop-down lists to make it easier for users to categorize events they are reporting according to the schema that Sahana uses.
- We need to develop the ability to push data from Sahana back to smart and simple cellphones. To accomplish this, we have multiple strategies: We plan to develop Android application that can display Sahana-plotted points on google maps app on the phone, and for future development, will do the same for iPhone and Blackberry over the next year. And we need to develop a means to push simple textual data to cellphones without a request needing to be sent (a subscription to certain categories of updates).

Do you want to attend the next RELIEF event: ABSOLUTELY. See you in February in Monterey!

Appendix 4: Synergy Strike Force

Under RFP 16019, Synergy Strike Force was asked to submit a requirements document and technical specification for the information sharing system that it has been developing at CP Roberts. These Documents are attached.

Operational Requirements Document

2009.11.21 (Draft 0.43)

Table of Contents

General Description of Operational Capability	35
Capability Gap.....	35
Overall Mission Area Description.....	36
Description of the Proposed System	37
Supporting Analysis.....	37
Mission the Proposed System Will Accomplish	38
Operational and Support Concept.....	38
<i>Concept of Operations</i>	38
<i>Support Concept</i>	39
Threat	41
Existing System Shortfalls	42
Capabilities Required	43
Operational Performance Parameters	43
<i>Function 1. Put/Get: Enable the user to deposit and withdraw information from an individual device.....</i>	<i>43</i>
<i>Function 2. Query. Enable the user to query/search the data stored on an individual device.</i>	<i>43</i>
<i>Function 3. Log. Enable each device to log all transactions on the system.</i>	<i>43</i>
<i>Function 4. Process/Store. Enable each device to process and store data deposited by a user.</i>	<i>43</i>
<i>Function 5. Send/Receive. Enable each device to transmit data from its storage subsystem and receive data from other devices.....</i>	<i>43</i>
Key Performance Parameters (KPPs).....	44
Simplicity	44
Openness.....	44
Extensibility.....	44
Portability	44
Durability	44
Scalability.....	44
Replicability	45
Wide Audience	45
Affordability.....	45
System Performance	45
<i>System Performance Parameters</i>	<i>45</i>
Interoperability.....	46
<i>Hardware</i>	<i>46</i>
<i>Software</i>	<i>46</i>
<i>Networking</i>	<i>47</i>
Human Interface Requirements	47
Logistics and Readiness.....	47
System Support	48
Maintenance.....	48
Supply	48
Support Equipment	48
Training	48

5.5 Transportation and Facilities	48
Schedule	49
System Affordability.....	50

General Description of Operational Capability

Capability Gap

In 2005, the Department of Defense (DoD) adopted stability operations as a core mission (See DoDD 3000.05). Stability operations were defined as

“an overarching term encompassing various military missions, tasks, and activities conducted outside the United States in coordination with other instruments of national power to maintain or reestablish a safe and secure environment, provide essential governmental services, emergency infrastructure reconstruction, and humanitarian relief.”

As part of this new responsibility, DoD specified that “integrated civilian and military efforts are essential to the conduct of successful stability operations,” and assigned responsibilities for collaborating with non-DoD entities to plan, prepare, and conduct stability operations. One of these responsibilities was for the DoD CIO to establish “policies and standards for technical information exchange and communications between the Department of Defense and other U.S. Government agencies, foreign governments and security forces, international organizations, and the private sector.”

In support of this new mission, DoDD 8220.02 (2009) clarified the information sharing requirements with stability operation partners. This Directive stipulated that:

"It is DoD policy that... to the extent authorized by law, and subject to applicable statutory and regulatory restrictions and limitations, information-sharing activities that facilitate coordination and cooperation between DoD and non-DoD partners will be established to enable common understanding of the stabilization and reconstruction, disaster relief, and humanitarian and civic assistance environment; and to support an integrated Whole-of-Government response capability."

In practice, such coordinated cooperative actions for information sharing have been bounded by several severe limitations:

1. **Closed Networks.** Most of DoD's ICT programs were designed to prevent access by outside parties by limiting access to the transport mechanism, the network. One of the architectural legacies of this design choice is an inability to grant access to these ICT resources (e.g., SIPR and NIPR) to uncleared civilians and foreign nationals working for non-DoD stability operations partners, because access to these resources would first require granting access to secure networks to these external partners.
2. **Differential Rates of Change between DoD and non-DoD Technologies.** The DoD tends to acquire custom-built, proprietary software through programs of record with long waterfall-based development and testing cycles. In

contrast, many non-government organizations and private sector firms acquire CoT and free/open source technologies, which tend to be released on quick, iterative cycles. During a crisis, it is common for a large pool of open source developers to release multiple versions of software in a single day in an effort to support humanitarian relief efforts. DoD tools are usually incapable of incorporating new channels, methods, or features on a timeline matched to such a tempo.

3. **Lack of Bandwidth.** In the field, bandwidth for communications is usually constrained or simply unavailable, either for technical reasons such as equipment failure; physical reasons such as weight or size restrictions on the equipment; or budget considerations (civilian access to satellite bandwidth is very expensive, often costing \$7 per megabyte). In this context, both DoD and non-DoD personnel are forced to resort to the physical exchange of memory devices (aka, 'sneakernet'). However, because DoD personnel cannot use the most common form of physical data exchange—the ubiquitous USB memory stick, and because both DoD and non-DoD personnel are often reluctant to open spinning drives (CD and floppies) to the sand and dust that is endemic to austere environments, information sharing is often one way: towards the DoD. Such a dynamic creates few incentives for non-DoD organizations to share data that often contains valuable HUMINT.

The combined effects of these limitations—closed military networks, quickly changing ICT systems, and low bandwidth—is the gradual insulation of the DoD from the information flows of its stability operation partners and the perceived withdrawal of DoD interest in working with those partners.

DoD requires an approach to bridging the gap between the mandated information sharing under DoDDs 3000.05 and 8220.02 and the observed information sharing available from existing tools and observed during operations in Indonesia (2004-05 Tsunami, Iraq, and Afghanistan).

Overall Mission Area Description

This ORD addresses a mission area from DoDD 3000.05 and DoDD 8220.02: creating the minimum essential capabilities required to enable information sharing between DoD and non-DoD actors involved in stability operations. This ORD outlines the requirements for an information exchange environment that will facilitate information sharing among the non-combat related organizations in a conflict prone region, with special focus on agriculture, public health, and education.

It is the intent of this effort to develop workable capabilities that enable unclassified information to be shared among the elements that are engaged in civilian-military operations, for the purpose of improving information flow on stability-related projects. The system here outlined shall support situation monitoring, information sharing, billboard, news services, status of needs and their resolution, and additional data

needs. It is anticipated that these new capabilities will be used to scale solutions to multiple locations.

Description of the Proposed System

To bridge the capability gap between the ICT systems of partners to stability operations, a solution must create an information-sharing environment that provides users with physical access to a device that can store, process, and make sense out of the partner's own operational data. This solution should be capable of working as a single standalone unit as well as a system of distributed devices connected by whatever networking channel is available to the user. The larger system of interconnected devices will form an information exchange system, where users can deposit and withdraw data about ongoing stability operations.

The system has five primary requirements:

1. **Simplicity.** The system shall be designed with the minimal essential functionality to meet its mission. Simplicity of design will drive lower costs of production, training, implementation, and maintenance. Simplicity will contribute to the usability of the system.
2. **Standalone Capability.** Each individual device in the system must provide minimal data processing, visualization, and data storage to enable users to make sense of their own data sets. Each device must translate support common data formats used by both DoD and non-DoD partners (including office productivity and mapping formats).
3. **Suitability for Field Environments.** Each individual device must be tooled for use in austere field conditions. It will support use in contexts where access to electrical power is intermittent and disconnected use is common.
4. **Scalability.** Individual devices must network into a distributed network of peer devices, enabling users to choose to exchange information with a network of other organizations using the device and/or a standard API for interacting with the devices.
5. **Extensibility.** To ensure that the overall system can adapt to changing ICT methods and channels, the system should implement open standards and use CoT or FLOSS tools wherever possible.

Supporting Analysis

These requirements were derived through a series of interactions with:

1. Subject matter experts that have participated in civilian-military operations overseas;
2. Information technologists that have employed commercial off-the-shelf capabilities to assist in the aggregation, fusion and exchange of unclassified information; and
3. Individuals currently involved in efforts that leverage existing in-theater information sharing capabilities in previous and ongoing civilian-military operations.

Mission the Proposed System Will Accomplish

The proposed system will create an approach to meeting the mandates for civilian-military information sharing articulated in DoDDs 8220.02 and 3000.05. This mandate requires the DoD to support “stability operations activities led by other U.S. Government departments or agencies, foreign governments and security forces, international governmental organizations, or when otherwise directed.” One responsibility assigned to the CIO of the DoD is “identifying and developing strategies for the use of ICT capabilities to enable civil-military interaction, information sharing, and accelerating stability and reconstruction activities.” This ORD directly meets this information-sharing mission.

Operational and Support Concept

Concept of Operations

The end users of the proposed information sharing system operate in a context in which the underlying assumptions of the developed world are inoperative. At best, they expect intermittent power and networking connectivity under primitive conditions. During stability operations, all users—military and civilian—work under high stress and information overload (in multiple languages) while tired, sweaty, dusty, and hungry.

However, unlike their counterparts in uniform, the staff of NGOs and other civilian organizations cannot assume reachback support, quick availability of reliable airlift or ground transport of essential supplies, or a large hierarchical organization which can provide intel, imagery, or decision support on a timeline that meets operational tempi. Staff at these civilian organizations tend to migrate from crisis to crisis, adapting their methods and toolsets to local circumstances. As a result, their organizations tend to be decentralized, with decision making authority delegated close to the ground worker. They also tend to rely less on reaching back to a distant office for advice; instead, they build networks of trust among the community of people who operate in the immediate vicinity (and who meet each other during responses to crises around the world). For them, crisis is not a carefully scripted musical score, but an improvisation more akin to jazz.

Their kit tends to reflect this sense of self-reliance, community, and adaptability. They operate with a mantra, “only pack it if you can hack it.” (Hack being defined in the traditional sense of “to make elegant modifications to” rather than the incorrect conflation of hack with crack, meaning “to break security of.”). That is, NGO staff tend to bring kit which enables customization, modification, and mashup; they eschew closed source, hermetically sealed tools where the software is burned into the unit (though an exception appears to exist for office productivity software, where all parties are often locked into proprietary software like Microsoft Office). NGO staff are required to constantly adapt, and expect that their tools will be flexible, scalable, and sufficiently easy-to-use that they can teach local nationals how to make effective use of them.

The processes required under traditional USG acquisition process is often mismatched to systems which must evolve in the field. Requirements documents ask vendors to project stable functionality and test to those functionality with metrics for performance set over a longer time horizon than the release cycles of open source software. In many cases, the use case of large programs of record is considered to be sufficiently stable to justify this type of rigid development. In the case of stability operations, however, such inflexibility is a design flaw.

Current systems used by NGOs include a medley of proprietary software (often Microsoft products) mixed with some web-based tools and custom applications, built either by proprietary vendors for specific use cases or home-grown tools without careful thought about interoperability or data exchange outside the organizations that built them. They tend to meet only local needs and store data in proprietary (binary) formats. This situation creates a well-known barrier to sharing information between organizations: the need to re-enter/rekey data into the target system. NGO staff barely have sufficient time to enter data in one system; multiple data entry is simply an impractical solution.

The proposed information sharing system addresses these issues by creating a platform for improvisation, built on a simple framework that is modifiable ("hackable"), flexible, easy-to-use, and built on open (interoperable) standards. The intent is not to build a traditional system, with fixed support desks and a proprietary network. Instead, the intent is to build knowledge of how to modify the system directly into the community and to make use of existing bandwidth. This system will not care what protocol data packets flow through; it will make use of existing pathways. It will read commonly available data formats and make them available to all parties. It will also use XML-based data feeds to limit the need to re-enter data (instead opening the possibility of an API/Web Service which can consume data from operational partners).

Support Concept

Because the end users of this system operate under intermittent connectivity, and because their bandwidth costs are often prohibitively expensive (BGAN costs \$7/MB), this system needs to replace traditional help desk support with redundant alternatives. This system uses two approaches to support:

1. The social model pioneered under open source software like Apache and Drupal. For these applications, there is no help desk; rather, users with problems or bugs submit their issues to community forum, and other users submit their insights, responses, and quite often, working code patches.
2. A network of experts who will operate in theatre, visiting the users of systems and bootstrapping the capabilities of NGO and military technologists to hack/mod the systems themselves. These experts will be trained by the vendor both in how to use every part of the system and to train others how to train their own staff.

Critically, the system will also create the same "stone soup" dynamic behind the success of OSD initiatives like Strong Angel disaster response demonstration series. By creating a modifiable platform, the individuals within the community can band together and can add functionality that meets their local needs while simultaneously building the overall capacity of the entire community involved in a stability operation. The vendor will study Strong Angel and provide the social organizing and platform for this stone-soup dynamic to occur.

In this light, it is vital to add that a minimum essential system is not intended to be a perfect system. As stated in the COIN manual, "The hosts doing something tolerably is often better than foreigners doing it well." Here, embedding the skills for NGOs and local nationals to operate a device tolerably is far more effective than a central organization running it well. The objective is to provide a system with minimum complexity and maximum flexibility, using familiar open standards and browser-based interfaces to enable NGO staff to operate and maintain the device themselves--tolerably and perhaps not perfectly, but the system is their own to modify and improve upon. They can check the code and ensure that the DoD has not inserted any backdoors, and they themselves can determine the level of trust that they place in the system. While this design may not allow control of the system, it does enable the user base to innovate and meet the requirement for flexibility. It also piggybacks on the extant communities of practice which emerge across NGOs, enabling technologists and other SMEs to rely on each other for support across organizational boundaries.

For organizations that require help desk support and are willing to pay for reach-back from the distant field by available communications channels, there exist commercial support contracts for many of the applications we are suggesting for use within the device (such as GeoCommons, Drupal, etc). Numerous system integrators exist can take offer support contracts on this design for a prototype information sharing system. Help desk support is beyond the scope of this prototype.

It is also important to note that the systems would be distributed with a set of initial data, populated with both applications and data appropriate to the local context. This initial seed--the stone in the soup--has repeatedly proven to be the essential catalyst for information sharing. By giving something away, users are willing to provide information in return. In our experience, the most valuable currency has been updated satellite imagery.

Threat

This ORD addresses a threat posed in the revised Counterinsurgency Manual: "Sometimes the more you protect your force, the less secure you may be."

Field experience has repeatedly shown that locking down information foments uncoordinated actions by NGOs and military units, ultimately breaking down unity of effort. Without unity of effort, the responding organizations cannot supply necessary services to local population, diminishing the legitimacy and efficacy of the overall operation.

Stability operations tend to be dynamic: after a major shock to an affected nation (such the diminution of combat activities or a natural disaster), military and civilian organizations scale their participation over time. During this increase in the scale of operations, information-sharing problems quickly lead to coordination problems: as more donors and projects emerge, the need for coordination increases at an increasing rate. However, actual coordination between organizations rarely keeps pace with the desired level of coordination. More commonly, information shared between stability operation partners diminishes over time, usually as concerns over force protection and desires to prevent insurgents from discovering and thwarting activities of stability operation partners creates increased focus on information assurance.

This focus on protecting information usually leads to conflicts and a concomitant breakdown in trust between the independent actors. Thus begins a vicious cycle: as trust decreases, the amount of information flowing between actors decreases, leading to further breakdown in coordination, more conflict, and ultimately yet more decreasing trust and reduced information sharing.

These challenges of information sharing are large factor in the downward spirals of cooperation and the opening for insurgent activities after a disaster or conflict. When trust dissipates between stability operation partners, information sharing gets limited to carefully prescribed reports and ground truth gets lost to all but those who are closest to the affected population (who may well be insurgents). Though the degradation of trust and coordination, stability operations turn into counterinsurgencies, making any civilian operations far more risky, costly, and ultimately, less effective.

Existing System Shortfalls

Systems for information sharing between organizations are largely absent in theatre; where found, current DoD systems exhibit the following problems:

1. **Closed to partners.** Networks are secure and therefore not open to use by non-DoD personnel.
2. **Not sharable.** Information exists under classifications or policy restrictions that prevents its use by other organizations.
3. **Not configured for portable austere use.** Many servers require constant Internet access and electricity, and are not configured for use in conditions with intermittent networking and power.
4. **Not configured to enable anonymous submissions.** Most systems are configured with authentication/authorization tied to individual identities. In cases where organizations are sharing information that might have sensitivities around its origin and they are uploading those data into a publically available archive, anonymous submissions are mandatory.
5. **Not sufficiently flexible to accommodate and adapt to change.** Programs of record require years to maneuver changes through development and testing and evaluation. Stability operation partners iterate technology at a far faster rate. Experience on the ground indicates that DoD technology for information sharing is falling behind stability operation partners at an increasing rate.

It is imperative to highlight the negative effects of overzealous information assurance policies on the ability of stability operations partners to achieve unity of effort. As field operations in Iraq and Afghanistan continue to reveal, when critical information on development activities (such as building a school or water treatment system) is withheld from stability operation partners, it rarely has the intended effect of keeping insurgents in the dark about the operation, and nearly always causes lack of coordination between DoD agencies, NGOs, and partnering local nationals. This lack of coordination causes duplication of effort, loss of credibility, and mission failure. New systems for information sharing should not replicate the current silos between the military and host-national and international partners.

Capabilities Required

Operational Performance Parameters

4.1.1. Users.

The system will support the concepts of users and groups. Anonymous users will be able to submit information, but that information will remain unverified. Verified users can login and share information which can be private to them as individuals, private to a group or groups, or public to the users of the system.

At a minimum, each device in the system shall be enable each user to:

Function 1. Put/Get. Enable the user to deposit and withdraw information from an individual device.

The system shall support standard data formats used by non-DoD entities, including common word processing, spreadsheet, presentation, database, mapping, imaging (still and video), and web services formats.

Function 2. Query. Enable the user to query/search the data stored on an individual device.

The system shall support queries across document formats.

4.1.2 Devices.

At a minimum, each individual device in the system shall be capable of the following functions:

Function 3. Log. Enable each device to log all transactions on the system.

Each device shall support reports of all transactions which tie into measures of effectiveness.

Function 4. Process/Store. Enable each device to process and store data deposited by a user.

Each device shall support extensible software for processing data, including methods to a) parse common formats for structured data; b) display structured data in charts, maps, and graphs; and c) display an index of unstructured data stored on the system.

4.1.3. Network.

At a minimum, the network that connects the system of devices shall be able to:

Function 5. Send/Receive. Enable each device to transmit data from its storage subsystem and receive data from other devices.

The system shall be able to use the best available networking connectivity, including the connections commonly found in austere conditions and networked use without connection to the public Internet.

Key Performance Parameters (KPPs)

Simplicity

The system shall be built to meet the minimum requirements for catalyzing information sharing within and between the DoD and non-DoD partners to stability operations. The hardware design shall be as simple as is needed for austere conditions. The software design shall be as simple as necessary to enable core user, system, and network functions outlined below.

Openness

The system shall remain as openly accessible as is reasonable in order to ensure ongoing utility. For the purposes of acceptance/rejection, open shall be defined as adhering to common open standards and data formats for all its included software, including but not limited to: vector GIS data (OpenStreetMap compatibility), raster GIS data (OpenLayers compatibility), structured data (XML and MySQL), unstructured data (OpenDocument (XML) and W3C XML Schema common the field like XHTML and CSS), and common productivity software (PDF, Microsoft Office).

Extensibility

The system shall be designed to have extensible hardware and software, with designs that enable quick integration of new tools and technologies. For the purposes of acceptance/rejection, extensible shall be defined as having an operating system and applications that is open source and can be checked in and out of version control.

Portability

The system shall be easily portable in the field. For the purposes of acceptance/rejection, portable shall be defined as having a form factor that can be carried by 1-2 persons between remote locations without risk of irreparable damage to the internal hardware.

Durability

The system shall be built to maximum its durability in the austere conditions of the field, balanced with simplicity and affordability. For the purposes of acceptance/rejection, durable shall be defined as capable of working in a covered primitive structure without unacceptable risk to the internal hardware from water or sand. The system shall not harden the case for MILSPEC, nor shall the system increase component costs to beyond 2 times the rate of commercial, off-the-shelf computing system of similar specifications (processing, memory, and storage).

Scalability

The distributed system of devices shall be scalable to the limits of a stability operation and within the limits of affordability and simplicity. For the purposes of acceptance/rejection, scalable shall be defined as enabling remote software updates to the operating system and accompanying applications; providing at least 1 free USB port for new hardware; providing replaceable storage device; and enabling one

system to send/receive data directly with another system via any of the supporting networking protocols.

Replicability

The system shall be easily replicable. For the purposes of acceptance/rejection, replicable shall be defined as having the ability to clone one system—including its entire operating system, applications, data, and current configurations—to another blank device within 24 hours.

Wide Audience

The system shall be designed to the widest possible audience within the limits of affordability and simplicity. For the purposes of acceptance/rejection, wide audience shall be defined as fielded UN agencies (e.g. UNICEF, UN HABITAT, etc), fielded NGOs (e.g., CARE, ActionAid, Worldvision, etc), and U.S. military units directly involved in supporting stability operations (e.g., Civil Affairs, HTTS, PDTs, ADTs, etc). The system—which is focused on stability operations—will not attempt to support the unique needs of combat units, though no barrier will be placed for their use by such units.

Affordability

The system shall be built from CoT and free/open source software to minimize costs, within the performance limits of durability, portability, and scalability. For the purposes of acceptance/rejection, affordability shall be defined as having a target unit price of less than \$5000 USD.

System Performance

System Performance Parameters

KPPs for Device Hardware

- Storage:** storage capacity is appropriate to medium-term persistence of imagery, video, and photography as well as databases and document repositories.
- Processing:** processing capacity is appropriate to fusing datasets. Processing will **not** be expected to be appropriate to video editing or imagery tile generation.
- Display:** provide minimum display necessary to support viewing of satellite imagery.
- Environmental:** Housing and hardware appropriate for use in both sandy environments and tropical environments, from winter cold to summer heat. To keep costs low and maintain simplicity, devices will not be hardened to MILSPEC, but will be expected to be used within shelters that protect them from water and dust to the maximum extent possible.

KPPs for Device Software

- Document Viewing:** software able to present usable views of common document formats, including word processing, spreadsheet, presentation

(slides), and PDF; common Internet-based applications, including Web, social media, and email; and common data visualization formats, including maps, charts, and graphs.

- Document Searching:** software able to search for keywords across document formats and within the database of structured data.
- Data Visualizations:** software able to create views of structured data in maps (where georeferences exist), charts, and graphs.
- Logging.** Able to log all transactions on the device. Able to allow both anonymous and non-anonymous submission of data.

KPPs for Users

- Able to upload/deposit common document formats to the device.
- Able to download/withdraw common documents to external media drives (e.g., user's own USB memory stick).
- Able to search device for relevant information by keyword.
- Able to visualize processed data in maps, charts, and graphs.
- Able to determine their own patterns of use.

KPPs for Administrators

- Able to view transaction logs.
- Able to manage access to the device.
- Able to extend software with patches and updates in the field.
- Able to replace and/or repair hardware in the field.

Interoperability

The device will use CoTs and Open Source tools that comply with standards for data exchange and open standards for documents, maps, messages, images, and videos. These standards will ensure current and future compliance with open standards for data exchange. For the purposes of acceptance/rejection, the device will support at the following data exchange formats: KML, XML, RSS, Email, and GeoRSS.

Hardware

- Power:** able to run on voltages worldwide (e.g., 100/240V). For the purposes of acceptance/rejection, the device will operate on voltages ranging from 100V to 240V.
- Ports:** at minimum offer commonly available ports. For the purposes of acceptance/rejection, the device shall include ports for USB2.0, RJ-45/Ethernet, and SVGA.

Software

- Documents:** Able to store, view, and share common office productivity data, mapping data, images, videos and support standard protocols for messaging, social media, and web browsing. For the purposes of acceptance/rejection, the device shall support the following document formats: KML, HTML, XML, RSS, GeoRSS, JPEG, MPEG, and TIFF

Networking

Device will be able to operate on best available networks, which may include cellular, satellite, WiFi, and wired Ethernet.

In the field, there a variety of networking protocols: BGAN, VSAT, WiMAX, WiFi shots, cellular data connections, and even TCP/IP over HAM radio. That said, these networking protocols generally interface with computing devices in only three ways: an RJ-45 jack (Ethernet), a USB port (for cellular data and some Wifi), and/or a Wifi radio antenna. For the purposes of acceptance/rejection, the device will support networking via an RJ-45 Jack, a USB port, and a WiFi antenna.

Human Interface Requirements

The system will operate through a standard web interface to minimize training requirements and technical skill. The system shall enable users to determine their own patterns of use.

Note: as a core requirement for building support among non-DoD users for this system, this requirements document will be opened to community review. This use of collective intelligence will not only guarantee that the design requirements for the device will meet needs from a large community of users, but that the final design will be familiar and perceived to be a community resource, subject to continuous process improvement.

Training in the analysis of data (and the associated knowledge domains) stored on the system is beyond the scope of this ORD.

Logistics and Readiness

The devices will be set up by trained facilitators, who will carry devices to and from sites where local partners are working on stability operations. Facilitators will have responsibility for administration of the hardware, software, and networking functions of each device in the distributed network of information-sharing systems.

System Support

Maintenance

The device must be field maintainable with minimum requirements for training. Devices will be maintained by facilitators who will be trained to replace parts which may wear out or fail, such as storage devices and networking equipment. Facilitators will also monitor drive space consumed and ensure that adequate storage space exists for continuous use.

Supply

The device will have an open design based on CoT and Open Source tools. All parts must be available CoT, enabling this project to leverage off existing supply chains for parts. Replacement parts should be ambiently available or source-able in most environments.

Support Equipment

The device will require a power and networking infrastructure. Software will be developed in standard integrated development environments and placed under revision control.

Training

Facilitators will require training in the configuration, operation, and replacement of hardware, software, and networking devices that compose the device. Users will require minimal training in how to upload/download files, search the device, and use the visualization tools.

Analysis of the data stored on the device will be the responsibility of the NGO field staff. This approach is in line with the paradox from the COIN doctrine: "The hosts doing something tolerably is often better than foreigners doing it well." NGO field staff need to learn to analyze data moderately well than to accept the long delays of centralized, DoD funded, SME-based analysis. That said, the device will support XML-based feeds, which can be ingested by high-end analysis tools and provided back to the participating NGOs and military units as processed, trusted analysis. The provision of such ongoing analyses is beyond the scope of the information sharing system.

5.5 Transportation and Facilities

Devices will be portable and available under an open design. Where necessary, facilitators will keep spare devices on hand for quick replacement of failed units.

Schedule

The design of the system is expected to occur using an iterative design method. An early prototype system will be rolled out in late November 2009. Based on experience gained with the early prototype system implementation, a CONOP will be released in late January 2010, followed by delivery of rollouts of production units over the course of 2010.

System Affordability

The system aims to be affordable, with an estimated per device cost of \$5000 USD. This cost does not recurring cost of Internet access.

Technical Requirements Document

2009.11.21 (Draft 0.43)

Table of Contents

1	Scope	53
1.1	Objectives	53
1.2	System Overview	53
1.3	Audiences and Use Cases	53
1.3.1	Audiences	53
1.3.2	Expected Use Cases (Examples)	53
1.4	Design Assumptions	56
1.4.1	Literate practitioners	56
1.4.2	Counterinsurgency Mindset	56
1.4.3	Open Architecture for Flexibility and Adaptability	56
1.4.4	Augment Existing Resources rather than Substitute for them	57
1.4.5	Bandwidth	57
1.4.6	Adequacy of CoT and FOSS tools to meet majority of use cases of austere conditions	57
1.4.7	Best Attempts to Adhere to Best Practices from Usability.gov	57
2	Applicable Documents	58
2.1	Government Documents	58
2.2	Commercial Standards	58
2.3	Open Source Standards	58
2.3.1	Protocols	58
2.3.2	Data Standards	59
3	System Requirements	60
3.1	Description	60
3.1.1	System Components	60
3.2	System Performance (Technical Requirements and Constraints)	61
3.2.1	Hardware	61
3.2.2	Virtual Machine Software	61
3.2.3	Application Software Package (ASP)	62
3.2.4	Interface Requirements	64
3.2.5	Physical Form Factor	64
3.3	Design	64
3.3.1	Hardware	64
3.3.2	Virtual Machine Client Software	65
3.3.3	Application Software Package	65
3.4	Documentation	66
4	Quality Assurance Provisions	67
4.1	General	67
4.1.1	Responsibility for Tests	67
4.2	Preliminary Acceptance Tests	67
4.3	Final Acceptance Tests	67
5	Preparation for Delivery	68

Scope

This TRD specifies the design of a Minimum Essential Software and Services for Knowledge and Information Transfer (MESS-KIT) to catalyze information sharing between civilian and military partners to stability and counterinsurgency operations.

Objectives

This system will improve the ability of partners to stability operations to share unclassified information about reconstruction activities.

This kit will focus on radical simplicity of design, encapsulating complexity wherever possible within modules of either Free and Open Source Software (FOSS) or Commercial, Off-the-Shelf (COTS) software. It will also harness social networks that already exist within and between organizations to accelerate adoption of the platform and to catalyze information exchange.

System Overview

This Minimum Essential Services and Software Kit (MESS-KIT) shall be composed of three (3) elements:

1. **APPLICATION SOFTWARE PACKAGE:** One or more Virtual Machine Instances that package together an operating system with a web server environment and all FOSS/COTS software modules. Example: A VMware instance of an Ubuntu Linux installation with a full LAMP web server hosting environment and associated web software.
2. **VIRTUAL MACHINE CLIENT SOFTWARE:** One Virtual Machine Software Client to package, distribute, and host one or more Application Software Packages and abstract the application software from the host operating system. Examples: VMWare Fusion and Sun VirtualBox.
3. **HARDWARE DEVICE.** A hardware device on which the Virtual Machine Client Software and Application Software Package will run. The Hardware will include a host operating system. Examples: MacMini running OSX, ASUS eeePC Netbook running eeeBuntu Linux.

Audiences and Use Cases

Audiences

The MESS-KIT will address the needs of a wide audience, focusing on non-governmental organizations (NGOs), private volunteer organizations (PVOs), local NGOs, local governments, provincial reconstruction teams (PRTs), agricultural development teams (ADTs), human terrain teams (HTTs), and other partners to stability operations.

Expected Use Cases (Examples)

The following list of use cases are not comprehensive nor prescriptive of use cases that the Contractor is required to support. Rather, they point to potential real-world

applications of the device and provide a sense of the flow of information between users.

Situational Awareness within a single partner organization

Bob, a member of a PVO uploads a spreadsheet of irrigation projects (with location data) to a document management system. Kris, a second member of the PVO, sees an RSS feed indicating that Bob has uploaded new spreadsheet about irrigation. Kris retrieves the document and inserts it into a mapping tool, which already has data on other projects (like roads and microhydro electric generation sites). The tool's geocoder plots the approximate locations of the irrigation project and puts the data onto a map of the region. Both Kris and Bob are able to view the map and discuss projects in context of other georeferenced data, including roads and proposed microhydro generation projects.

Situational Awareness between Intermittently Connected Partners

Razzaq, a logistics officer for a big NGO, manages the supply chain leading into a province experiencing mass human migration around a military operation. He is connected to his NGO's information systems by a VSAT, which he uses to communicate his calculations of supplies of water and food to his international HQ.

Floyd leads a PVO team that monitors human migrations. Floyd heard of the MESS-KIT Floyd received permission from his manager to share limited, reviewed information about the refugee situation via his MESS-KIT with Razzaq. Floyd uploads images from his COTS camera with the grid coordinates of new areas where refugees are congregating in the woods.

Razzaq's MESS-KIT then receives an RSS feed indicating new information has arrived from Floyd. Razzaq reviews the information, plotting the new locations of refugees and examining the apparent health of people in Floyd's photos. Razzaq begins planning how to securely investigate the situation with refugees who are hiding in the woods several days earlier than he would have otherwise been able.

Cross-Organizational Project Planning and Coordination (Act)

Drew, a facilitator for a World Bank Development Project, manages the mobilization of five villages for a post-earthquake, block-grant development project. Drew is often in the field and disconnected more than 80% of the time. He cooperates closely with Sheila, a USAID employee. Sheila herself is in the field more than 25% of the time. Both use separate instances of the MESS-KIT: Drew's on a laptop that he carries into the field, and Sheila's installed on a desktop in her unit's headquarters. Sheila has configured her system to pull RSS feeds from Drew's laptop when it is connected to the local network (not a public internet)

Drew creates a blog post in the MESS-KIT after each village mobilization meeting as a means of keeping minutes. He also records which village elders attended each meeting and notes the name, type, and location of all proposed projects under the

block-grant program in two spreadsheets (attendance.xls and proposed projects.xls). Drew uploads the two spreadsheets to the document management system.

When Drew gets to a network connection, an RSS feed tells Sheila about all Drew's blog posts as well as the creation of the two spreadsheets. Sheila passes the attendance spreadsheet onto her colleague, Travis, who updates a sociogram, and she uploads the spreadsheet of proposed projects into her MESS-KIT's mapping tool. The tool's geocoder plots the approximate locations of the projects and puts the data onto a map of the region. Sheila compares these projects against plots of other proposed projects from NGOs in the region, and notes that the one NGO has already received funding to build a health clinic in one of Drew's villages which had decided to build one of their one. Sheila sends an email to Drew noting the possible conflict.

Extending the Information Sharing Platform

Craig is a fielded IT staff member of the UNJLC who has received a MESS-KIT as part of a decision by his Health Cluster to deploy and support the technology. Early in a deployment to a post-conflict situation, he notices a sharp uptick in the number of field assessment forms that are arriving with hard-written notes in the margin that state: "noted (n) persons with signs of mutilations by gangs of youth." Craig realizes that partners to the stability operation need to monitor this emerging situation and quantify the scale of this new problem. He modifies the data schema of the disaster management system, adds code to support a new set of fields about gangs and mutilation, and creates a patch which other IT staff members can install on their MESS-KITs. He uploads the patch to the document management system and emails his peers about it. His peers download the patch via a link in their RSS feed readers, and test the patch on a local non-production version of their MESS-KITs. Several submit improvements and a bug fix to the patch. Within several hours, the patch is ready for everyone to install on their production MESS-KITs.

IT staff across the stability operation install the patch, modifying their own disaster management systems and associates field assessment forms. The next day, more than 80% of the field assessment teams are taking quantifiable measurements about gangs and mutilation.

Maintaining Systems

Dr. Ashahi, an infectious disease specialist consulting to a host nation's ministry of public health, is responsible for improving an avian flu health project in a post-conflict region. She has a MESS-KIT in her field office on a MacMini, which she uses to track locations of outbreaks, note operational data about clinics, and receive outbreak data and maps from adjacent regions. Due to nearly 24/7 overuse, her hard drive crashes. She has an external hard drive, where she has stored data up to the previous day by cloning the virtual machine each night to the disk. She has been trained in what that she can open her laptop, where she has installed a spare copy of the Virtual Machine Client Software. She connects the external hard drive to her laptop, opens the last-saved version of the MESS-KIT virtual machine, and continues from where she left off the night before. She re-enters data from the current day, notifies the ministry's IT

department that she needs a new hard drive, and closes the day by saving the virtual machine to the external hard drive and her local notebook drive.

A week later, an IT staff member from the ministry arrives with a new hard drive. He installs the drive on the MacMini. Ashahi then suspends her MESS-KIT's virtual machine, connects the external hard drive to the MacMini, copies the current state of the virtual machine onto the MacMini, and continues right where she left off on the MacMini.

Design Assumptions

Literate practitioners

This design assumes reading literacy and computer literacy on the part of users as a minimum condition for use. It also assumes literacy with the use of web browsers to access and upload data. The design will harness a growing pool of intelligent technologists in partner organizations, who are familiar with open source tools and share solutions in the field.

Counterinsurgency Mindset

This design assumes support for the application of the revised counterinsurgency (COIN) doctrine to the problems of information sharing during COIN/stability operations. In particular, it assumes both the assumption of increased risk and an understanding of a paradox from the COIN manual:

1-149. Sometimes the more you protect your forces, the less secure you may be.

Restated in IT terms, this paradox yields important insights into the nature of the MESS-KIT. Paradox 1-149 can be restated: "Sometimes the more an organization protects its information, the less secure it may be." This is particularly true when information assurance policies all but stop the flow of information about reconstruction activities both within and between operational partners. When these information silos ossify into inviolable domains, it is impossible to obtain COIN's Unity of Effort and insurgents will understand that they can take the strategic position of waiting out and wearing down opponents to their ends.

The MESS-KIT assumes a mindset which seeks to reverse this dynamic by opening flows of information between COIN/SSTR partners and concomitant efforts to improve Unity of Effort. In return, users of the system must assume additional risk related to their operations as well as the security of information.

Open Architecture for Flexibility and Adaptability.

Operations around COIN, HADR, and SSTR require constant adaptation and flexibility. No two operations are exactly the same, nor can a tactic in one locality or time period necessarily be applied effectively in another localities or time periods. As a result, the MESS-KIT assumes a mindset that values flexibility and adaptability to an extent that users will assume additional risks and eschew canned IT systems that may not be fully appropriate in the context or effective at confronting problems in that context. The

MESS-KIT assumes that the users will modify the software to meet current challenges, and will also share their adaptations with partners when appropriate.

Augment Existing Resources rather than Substitute for them.

In the field, new tools are not always welcomed, particularly when introduced during a period of stress. It is therefore assumed that the MESS-KIT will integrate with existing tools and create methods of making the best possible information flows given current constraints. The mosaic of IT systems will not be perfectly designed, but will catalyze gradual change towards common formats as the use cases for sharing become more clear to the users.

Bandwidth.

The MESS-KIT assumes availability of some bandwidth via common methods (BGAN, rBGAN, WiMAX to WiFi, WiFi, VSAT, POTS, or cellular data), or provision of bandwidth to stability operation partners under DoDD 8220.02 and 3000.05.

Adequacy of CoT and FOSS tools to meet majority of use cases of austere conditions.

The MESS-KIT assumes that FOSS and COTS hardware and software are adequate to field conditions and field needs of COIN, SSTR, and HADR operations. It assumes that mapping tools will be affordable for use in the system and/or available through free and open source tools (FOSS).

Best Attempts to Adhere to Best Practices from Usability.gov

The design will make best attempts to follow the guidelines established at Usability.gov. See <http://usability.gov/basics/>.

Applicable Documents

Government Documents

RFP 16019

ORD for RFP 16019

DoDD 8220.02

DoDD 3000.05

Commercial Standards

WiFi Networking (802.11b)

WiFi Networking (802.11g)

WiFi Networking (802.11n)

WiMAX

BGAN

rBGAN

VSAT

Open Source Standards

Protocols

LDAP Lightweight Directory Access Protocol

HTTP Hypertext Transfer Protocol

HTTPS Hypertext Transfer Protocol - Secure

FTP File Transfer Protocol

SFTP Secure File Transfer Protocol

SSH Secure Shell (<http://tools.ietf.org/html/rfc4251>)

SCP Secure Copy

WMS Web Mapping Service (<http://www.opengeospatial.org/standards/wms>)

IRC Internet Relay Chat (<http://www.irc.org/>)

XMPP Extensible Messaging and Presence Protocol (<http://xmpp.org/>)

Data Standards

XML	Extensible Markup Language
XHTML	Extensible Hypertext Markup Language
RSS	Really Simple Syndication
GeoRSS	Geo Really Simple Syndication
KML	Keyhole Markup Language
JPEG	Joint Photographic Experts Group (image format)
MPEG	Moving Picture Experts Group (video format)
REST	Representational State Transfer

System Requirements

Description

The MESS-KIT is intended to catalyze information sharing across civilian and military partners to SSTR, HADR, and Counterinsurgency operations.

System Components

Applications Software Package.

The ASP will perform five data functions:

1. **PUT/GET.** Enable the user to deposit and withdraw information from an individual device.
2. **QUERY.** Enable the user to query/search the data stored on an individual device.
3. **LOG.** Enable each device to log all transactions on the system.
4. **PROCESS/STORE.** Enable each device to process and store data deposited by a user. Note: This TRD will specify the methods for processing and storing data at the system level, but will not specify the methods used by individual FOSS/CoTS modules, which are beyond the control of the vendor.
5. **SEND/RECEIVE.** Enable each device to transmit data from its storage subsystem and receive data from other devices.

Virtual Machine Software Client

The containing Virtual Machine Software Client will perform three functions:

1. **PACKAGE.** Encapsulate a full MESS-Kit into a virtual machine that can be opened across platforms: Linux, Mac, and Windows.
2. **DISTRIBUTE.** Enable quick distribution and updating of the system via exchange of instances of virtual machines.

Hardware

The containing Hardware will perform four functions:

1. **STORE.** Keep a host operating system and virtual machine software on disk.
2. **HOST.** Host the operating system in which the virtual machine software will run.
3. **CONNECT.** Create a physical connection to networks and attached devices, internally and externally (storage, printing, keyboards, mice, displays).
4. **CARRY.** Encapsulate the entire system in a portable form factor that can be carried between locations.

NOTE: This TRD specifies the minimum hardware requirements to host an initial prototype of the VM + MESS-Kit. This hardware will host the environment wherein the virtual machine and MESS-Kit will reside. Over time, the highly portable nature of virtual machines will enable the MESS-Kit to reside on multiple platforms, from netbooks to data-center-level servers. This portability is an essential component of the scalability of the system.

System Performance (Technical Requirements and Constraints)

Hardware

Storage Space

Hardware shall provide storage capacity that is appropriate to medium-term persistence of imagery, video, and photography as well as databases and document repositories, no less than 200GB.

Processing Capacity

Hardware shall provide processing capacity that appropriate to fusing datasets, visualizing maps quickly, and hosting a web server environment for up to 12 simultaneous users. Processing will **not** be expected to be appropriate to video editing or imagery tile generation.

Latency

Hardware shall provide sufficient processing and memory capacity to permit average load times for web pages loaded from entirely local data via the web server environment to not exceed 10 seconds.

Graphics Capacity

Hardware shall provide minimum display necessary to support viewing of satellite imagery at a viewable resolution, no less than VGA (800 pixels by 600 pixels).

Operating System Support

Hardware shall support modern operating systems, including Ubuntu Linux, Apple OSX v 10.4, and/or Windows XP Pro.

Networking

Hardware shall support best available networking. Hardware shall support TCP/IP using PPP POTS modem, BGAN, VSAT, Wifi, cellular modem in the USB 2.0 port, or Ethernet CAT-5 cable in the RJ-45 connector.

Virtual Machine Software

VM Creation

VM Client Software shall provide the capacity to load current state of at least one operating system into a virtual machine.

Operating System Support

VM Client Software shall support at least one free and open source operating system which can be freely redistributed without royalty fees or licenses.

Packaging and Distribution

VM Client Software shall enable users to save the current state of an operating system and all current applications and data into a clone virtual machine. This clone can be opened on a second piece of Hardware (compliant with specifications in this TRD) and used with a copy of the VM Software Client.

VM Updates

VM Client Software shall provide the ability to download updates to the VM Client Software application via both Internet download and locally connected storage device (e.g., USB memory stick or external hard drive).

Application Software Package (ASP)

Simplicity. Small primitives.

ASP shall provide a set of minimum essential functionality for civilian-military information sharing, including the following functions, which may be drawn from either FOSS software, COTS software, or a combination of both (with preference give to FOSS software):

Access Control

ASP shall provide a method for controlling and logging access to the ASP (Log). This method shall provide access to individual files and applications by user and group. ASP shall default to setting information access rights to "global, public." ASP shall provide each user with the option of limiting access to a document to one or more groups, or to one or more users.

Data Backup

ASP shall provide a tool to backup data stored in associated databases and filesystems to an external device.

Web Server Environment

ASP shall provide a web server environment, including the web server software, database, programming languages, and configuration files necessary to run all hosted application installed by the contractor.

Web Browser

ASP shall provide a web browser that is aspires to be compliant with XHTML1.1 and CSS2 (note: no browser is 100% compliant).

Search

ASP shall provide a search capability (Query) which can perform a full-text search by keyword across the filesystem. Each application in the ASP shall provide a search capability which can perform a keyword search within its database.

Document Management System with Viewer

ASP shall provide a document management system where a user can Put/Get data: upload one or more files, categorize those files by a system of terms (taxonomy), download documents to a local environment or storage device, and view files in common file formats for stability operations (including but not limited to word processing, spreadsheet, images, video, maps, RSS feeds, and web pages).

Wiki

ASP shall provide a wiki tool, where users shall be able to Put/Get data: enter data in wiki format collaboratively; search entries in the wiki by keyword; be able to track revisions of each wiki page; be able to restore web pages to a previous state; and be able to discuss conflicts over each wiki page in a compatible discussion forum.

Blog

ASP shall provide a blogging tool, where users can Put/Get data: post short text entries, categorize those entries with terms, search across the blog, and browse previous entries by date and topic.

Image Gallery

ASP shall provide an image gallery, where users can Put/Get data: upload images from COTS cameras, categorize those images with tags, and view images in a gallery format.

Portal with Feed Reader

ASP shall provide a basic portal where XML feeds from each ASP function can be aggregated and viewed (Send/Receive).

Mapping Tools

ASP will attempt to provide an affordable mapping application which enables users to Put/Get data: upload georeferenced data into a geocoder and to plot georeferenced data on a basemap of the area. Note: Open-source mapping applications are still in nascent form; COTS application may be a required choice to support this requirement. If costs of these COTS tools escalated beyond 20% of the projected costs of the system, mapping tools may be dropped from the TRD.

Disaster Management System

ASP shall provide a disaster management tool which enables users to Put/Get data, including the ability to upload and view data about volunteers, upload, search and view data on missing people.

Interoperability

ASP shall provide a suite of applications which can expose lists of major entities in their data models to external applications via XML-based feeds (e.g., RSS, GeoRSS) and/or XMI-based APIs (e.g., REST).

Open Source

ASP shall provide applications whose source code is viewable and modifiable under common FOSS licenses, including GPL, LGPL, Creative Commons, BSD, and/or MIT.

System Administration

ASP shall provide a method for logging transactions on the system, including error messages, access control messages, and logs related to final Measures of Effectiveness (TBD).

Interface Requirements

Hardware Interfaces

Hardware shall provide at least one (1) USB 2.0 port, one (1) SVGA port (or DVI port with adapter), and one (1) RJ-45 Ethernet port.

Virtual Machine Interfaces

VM Client Software shall provide virtual interfaces to at least one (1) USB 2.0 port, one (1) video interface, and one (1) Ethernet port.

Application Software Package (ASP) Interfaces

ASP shall provide interfaces to Send/Receive data with peer ASP instances (local or remote) via two methods:

1. Clones of the virtual machine instance that contain the ASP.
2. XML data feeds to external applications (e.g. RSS, GeoRSS).

Physical Form Factor

Hardware

Portability

Hardware shall be portable by two persons (less than 40kg).

Environmental Factors

Hardware shall be appropriate for use in both sandy environments and tropical environments, from winter cold to summer heat. To keep costs low and maintain simplicity, Hardware will not be hardened to MILSPEC, but will be expected to be used within shelters that protect them from water and dust to the maximum extent possible.

Design

Hardware

Selection

Hardware will be selected based on a balance between affordability, portability, durability, scalability, and simplicity. No hardware shall cost more than 2 times the amount for equivalent functionality in a COTS device. Hardware selection will prefer COTS devices over custom-built solutions.

Extensibility

Hardware will be extensible and scalable. Hardware shall offer the interfaces specified in 3.2.4.1 and shall support the addition of external hardware devices, networking devices, and graphic devices via these interfaces.

Durability

Hardware shall be built to maximize its durability in the austere conditions of the field, balanced with simplicity and affordability. Hardware shall be capable of working in a covered primitive structure without unacceptable risk to the internal hardware from water or sand. Hardware shall not be expected to comply with MILSPEC, nor shall the system increase component costs to beyond 2 times the rate of commercial, off-the-shelf computing system of similar specifications (processing, memory, and storage).

Affordability

Hardware shall be affordable, both as a system as individual components for spares.

Security

Hardware shall not be expected to offer hardware-level security.

Virtual Machine Client Software

Selection

VM Client Software shall be selected based on its ability to meet the specifications in 3.2.2.

Portability

VM Client Software shall be able to generate virtual machines that are portable across MESS-KITs.

Scalability

VM Client Software shall be able to scale for use in mini-servers that are capable of providing ASP services to at least 12 people.

Affordability

VM Client Software shall be less than 10% of the system cost to license.

Security

VM Client Software shall not offer security beyond the access control to the operating system contained in a virtual machine instance.

Application Software Package

Selection

Applications selected for ASP shall support the functions listed in 3.2.3.1.

Extensibility and Openness

Applications selected for ASP shall be open source, defined as having source code that can be checked in and out of version control.

Simplicity

Applications for ASP shall provide minimum functionality. As a whole, ASP shall provide the basic functions for stability operations in compact, simple, maintainable package.

Scalability

Applications selected for ASP shall be scalable to at least 12 simultaneous users.

Affordability

Application selected for ASP shall not cost (individually) more than 20% of the total cost of the device.

Replicability

Applications selected for the ASP shall have configurations which are easily replicable across MESS-KITs.

Security

Applications selected for ASP shall support security by users, groups, and public access.

Log Message Generation

Applications selected for ASP shall provide one or more methods to log messages about errors, attempts to violate access controls, and current version of the software.

Documentation

Documentation shall be generated and delivered in accordance with the Statement of Work.

Quality Assurance Provisions

General

Prior to final Government acceptance, the MESS-Kit will be tested to ensure ... Testing shall be done to determine compliance with the requirements in Section 3 of this document.

Responsibility for Tests

The Contractor shall be responsible for verifying compliance with this specification. Government representatives (designated by the Sponsor) will witness testing at the Prototype Facility described in the Statement of Work and will certify results.

Preliminary Acceptance Tests

The Contractor shall be responsible verifying that the Prototype meets the requirements of this specification.

Final Acceptance Tests

The contractor shall conduct Final Acceptance Tests to verify the MESS-KIT's ability to support the requirements in this document.

Preparation for Delivery

The MESS-KIT shall be packaged for delivery according to best commercial practices and prepared for austere field conditions.


Supporting Documentation

2009.11.21 (Draft 0.43)


CONSEQUENCES OF POOR INFORMATION SHARING

Helping the Insurgency

1. DUPLICATE EFFORTS



Defense ICTs prevent access by outside parties by limiting access to the transport mechanism, the network. SIPR and NIPR are—by design—segregated from the public Internet used by stability operation partners.



Stability operation partners cannot get access to DoD ICT resources (e.g., anything on SIPR and NIPR), because network access cannot be granted to uncleared civilians and foreign nationals working for non-DoD stability operations partners.

2. BROKEN RELATIONSHIPS




DoD personnel cannot use the most common form of physical data exchange—the ubiquitous USB memory stick. Because both DoD and non-DoD personnel are often reluctant to open spinning drives (CD and floppies) to the sand and dust that is endemic to austere environments, information sharing is often one way: towards the DoD.




Stability operation partners often pay \$6/MB for satellite bandwidth to the public Internet from austere locations. They ration connectivity and find creative ways to use closed local subnets that never travel over the public internet or are accessible to DoD personnel.

3. LOST OPPORTUNITIES



The DoD tends to acquire custom-built, proprietary software through programs of record with long waterfall-based development and testing cycles. They are based on proprietary standards, closed data formats, and closed data schema that are inaccessible to NGO IT staff.



NGOs and private sector firms acquire CoT and free/open source technologies, which tend to be released on quick, iterative cycles. DoD tools are usually incapable of incorporating new channels, methods, or features on a timeline matched to such a tempo.


INFORMATION SHARING SYSTEM

Improving Unity of Effort through Open Data

1. OPEN ACCESS



Standalone Access. Stability partners on both sides of the civ-mil divide have physical access to a common platform for sharing operational information, which can store, process, and visualize each partner's own operational data.



Network Access. Partners using the devices collectively form a larger system of interconnected devices, creating an information exchange system, where users can deposit and withdraw data about ongoing stability operations.



Data-level Security. Devices individually and collectively use an open network with POSIX file-level security: users, groups, and public access. Anonymous access is supported as a core functionality.

2. BEST AVAILABLE NETWORKS



Wide Range of Networks. Systems will harness the best of the available networks, from WiFi to BGAN, VSAT, POTS modems, and cellular data.



Minimizing Bandwidth Use. The systems minimize throughput by exchanging XML feeds that list available data instead of sending complete files. Complete files are available on request.

3. ADAPTABLE FRAMEWORKS OF ESSENTIAL FUNCTIONS



Minimal Essential Software. Users will have access to a shared set of extensible core functions accessible via web browser: document management, blog, feed reader, wikis, disaster management, image gallery, and search.



Open Source and Commercial Tools. Applications will be supported by existing open-source software communities and commercial vendors. Applications can be modified to suit local context and operations.



Simplified Deployment via Virtual Machines. Virtualization removes layers of complexity around configuration and updating by encapsulating complete systems in portable virtual machines. VMs can be exchanged on the fly and can be used as backup systems on partner computers.

INFORMATION SHARING SYSTEM High-Level Architecture

