

# Sierra: Cooperative Request-Response for Resource Management in Disasters using Semantic Web Principles

Seyyed M. Shah, Christopher Brewster and Duncan Shaw

Aston University, Birmingham. B4 7ET  
{s.m.a.shah|c.a.brewster}@aston.ac.uk

**Abstract.** Disasters cause widespread harm and disrupt the normal functioning of society, and effective management requires the participation and cooperation of many actors. While advances in information and networking technology have made transmission of data easier than it ever has been before, communication and coordination of activities between actors remain exceptionally difficult. This paper employs semantic web technology and Linked Data principles to create a network of inter-communicating and inter-dependent on-line sites for managing resources. Each site publishes available resources openly and a lightweight open-data protocol is used to request and respond to requests for resources between sites in the network.

## 1 Introduction

A disaster<sup>1</sup> is defined by the UN as a “serious disruption of the functioning of a community or a society involving widespread human, material, economic or environmental losses and impacts, which exceeds the ability of the affected community or society to cope using its own resources” [1]. Disaster and emergency response is the most critical phase when emergency management agencies attempt to provide services and assistance so as to save lives, reduce the impacts on health, make sure that basic security and subsistence needs are met. In attempting to achieve this, disaster response agencies face a number of challenges. A number of different resources need to be identified, including people, physical resources, vehicles, etc. Agencies need to know their quantity, location with respect to the disaster and the time when they are available. Time and location are critical factors – resources provided too early, or too late or in the wrong place can reduce the effectiveness of recovery [2]. As well as time and location, collaboration is a critical factor. Usually any disaster of significant size will involve efforts and resources from many different and diverse organisations and individuals. Agencies need to know about these resources in detail so as to deploy them appropriately. Effective management of these issues fundamentally

---

<sup>1</sup> We use *disaster* and *emergency* interchangeably here although the UN defines an emergency as a “threatening condition” significantly less serious than a disaster.

concerns the availability and access to information about the resources. A given decision maker needs to know what, where, when and how for each and every resource whether a human being helping out or a load of sandbags needed to shore up a dyke in a flood.

This paper presents the prototype of a system built using semantic web technologies and based on Linked Data principles which seeks to address the information management challenges in disaster response. The paper is organised as follows: In Section 2, challenges in creating software systems for emergency response are discussed and set as evaluation criteria. Section 3 presents the cooperative architecture for decentralised resource management in emergencies, applying Linked Data principles for interoperability. A prototype implementation that demonstrates the the architecture is described in Section 4 along with the ontologies created for the implementation. Section 5 presents a Case Study that is used as an initial evaluation of the presented architecture.

## 2 Complexities of Software Tools and Resource Management in Emergency

Turoff et al. [3] discuss developing practical emergency response systems and present nine premises for preparedness, summarised here as follows:

- Regular Use: A system that is not used before a crisis will not be used during the crisis.
- Information Focus: Responders only use a system if the information provided is relevant to the emergency.
- Crisis Memory: Learning from previous crises is extremely important.
- Exceptions as Norms: Crises are unpredictable.
- Crisis Scope: Focus of problem solving shifts depending on the crises, each problem requires a range of different resources.
- Role Transferability: Responder roles can interchange during an emergency.
- Information Validity: Providing timely information is critical to decision making
- Open Information Exchange: Numerous responder organisations and individuals should exchange data freely.
- Unpredictable Coordination: Collaboration of responders and organisations is not pre-determinable.

For reasons of space, we will not analyse these premises further, but we will apply each premise to evaluate the work in this paper, and discuss how these goals might be achieved in practice. Existing disaster and emergency response systems used by emergency management agencies (EMAs), as our research has found, tend to be stand alone systems. Thus the fire service are unable to communicate electronically with the police and the police cannot communicate with the ambulance services etc. let alone with external third party data sources. Even within a specific service, there can be minimal data sharing between different offices and everything is communicated via faxes or emails.

This phenomenon of total lack of interoperability is also apparent in the software tools that have been developed to support humanitarian emergency response. In recent times, open-source online emergency response tools have been made available [4], and these have become widely used by NGOs. One of the most popular is Ushahidi [5] which allows crowd sourcing of information about an emergency. Any organisation can deploy Ushahidi to collect information from multiple “streams” for example, via text messages, email, twitter and web-forms, so that help can be sent where needed. Each deployment of the tool is completely independent of any other. This means that many organisations can simultaneously deploy the tool. However, due to the *centralised* nature of the tool, each organisation can only access information submitted to their own site. Furthermore, separate systems that gather emergency information from the public are not interoperable with Ushahidi. That is, the information collected about the same emergency in Ushahidi and a separate legacy system is not available to both sets of responders, even if there is overlap in the information gathered. Similar comments can be made concerning Sahana Eden [6] another popular humanitarian resource management tool.

This section has outlined some of the challenges for developing systems that are used in emergencies. The premises are set as evaluation criteria tools used in emergency that are centralised and not interoperable have been found. Following the next section, a system is proposed for decentralised and interoperable resource management in emergency scenarios, using semantic web technology.

## 2.1 Semantic Web Technologies

Semantic Web and Open Data technology builds on existing Web standards, RDF and OWL are standards used to add meaning to published data. Semantic technologies expect that data will be multi-authored and distributed, and crucially, interlinked between separate data sets. Semantic technologies also allows for creation and extension of formal ontologies, to provide a precise representations of aspects of the domain. Data from diverse systems can then be integrated and the ontologies to serve as a common language for communication between different organisations. Furthermore, separate systems using semantic standards can be integrated to create new applications for existing data.

Semantic technologies are therefore particularly suited to use in disaster management. Emergency response is typically multi-organisation and distributed. Using semantic technology has an important advantage- data from each organisation about the same subject, in this case an emergency, is linked together. Formal ontologies provide a common language for communication- both human and machine-based - about the emergency between diverse organisations. The graph-like nature of RDF allows navigation around data points and automated logical reasoning to improve the response. For example, the universal response database described by Turoff et al. [3], could be encoded using RDF and published “open data”. The next section describes a particular application of Semantic Web and Open Data technologies towards improving disaster response.

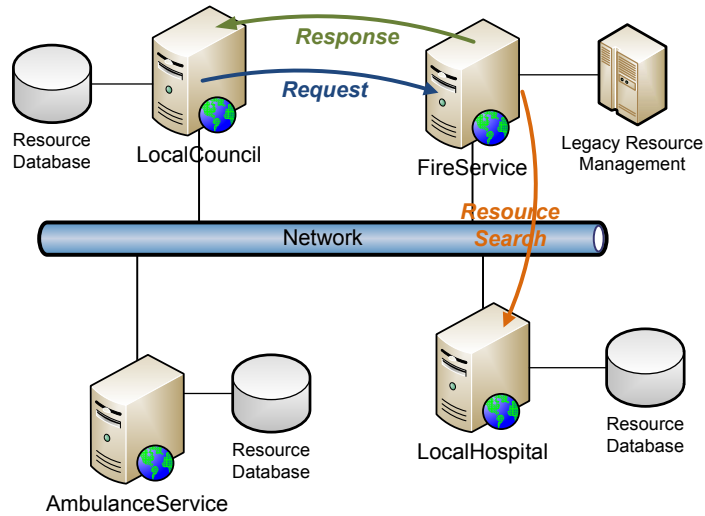


Fig. 1. Network of sites with resource responses and requests

### 3 Sierra Architecture for Decentralised and Interoperable Resource Management in Emergencies

In the light of the requirements identified in the previous section, we have developed a system based on the use of Semantic Technologies [7] and Linked Data principles [8,9]. An approach based on Semantic Technologies enables a decentralised and yet interoperable system design, a system which is both open to further expansion and yet scalable for complex resource management.

The proposed system is made up of a network of machine-readable websites. The purpose of each site is to hold information on resources, as well as act as an interface to make requests and respond to requests from other sites. Each site has a clearly defined remit e.g. the geographical area or organisation to which the resources, requests and responses belong. The basic architecture of such a network is shown in Figure 1.

Users of the site have access to the site (or sites) to which they are affiliated and a user can belong to multiple sites. The site presents tools that allow a user to view, modify and add data about resources concerning their particular organisation or entity. The web site also allows users to respond to requests for resources from any other site as well as make requests to other sites for resources in the network. User can only manage resources and make responses and requests from the sites to which they belong.

The users access the site via the Internet and a standard web browser. No additional software is needed, as the sites must use web standards for the user interface presentation. Using standards makes the sites accessible not only for a variety of devices - mobile, desktop or dedicated app, but also to disabled users

via screen readers and magnification. A web based system is key in this scenario because of the need for online collaboration in emergencies.

For interaction between sites, the network uses core principles of Linked Data. Firstly, data about physical resources that can be stored and exchanged between sites are encoded in RDF using an extensible ontology. This allows sites to exchange information about physical resources using a common, standards-based notation. Secondly, a lightweight semaphore protocol and closely related ontology is used to describe requests and responses for the exchange of physical resources. Each site runs a SPARQL endpoints so that remote sites can query them for information on resources, requests and responses.

A key aim of the system is decentralisation of sites that cooperate in the network. Each site and its resources are under the separate control of the organisation or group to which they belongs. As sites are web-based, sites can be physically located in geographically-diverse data centres. Web sites can be backed-up, mirrored and relocated as needed during an emergency.

Importantly, each site can be completely independently implemented, yet still be interoperable with other sites in the network. That is, if the protocol and standards are implemented correctly many different deployments can work together in a network. Interoperability via standards allows legacy systems to be integrated, for example the resource tracking systems used by some emergency responders could be integrated into the system with the addition of a compatibility layer.

## 4 Prototype Implementation

This section describes a light-weight proof of concept implementation based on the Drupal open source Content Management System (CMS). Drupal has been chosen because apart from being open source and extensible, it has a substantial set of contributed modules implementing semantic technologies and markup. Particularly important are the modules that allow for Drupal to become an open linked data system: RDF Extension, SPARQL, SPARQL views and Restful Web Services <sup>2</sup>. With the open data modules, content is manually mapped to a selected ontology, presented in RDF and queried using a SPARQL endpoint, via the ARC2 PHP library. Other contributed modules used included ones that enable geo-location: OpenLayers, geofield and geophp <sup>3</sup>. These modules collectively allow resources to have a location and to visualised using familiar online mapping interfaces such as Google maps or Open Street Map.

The prototype also customises Drupal by adding custom content types, so instead of managing blog posts and web pages, relevant kinds of content are managed. The content types added are *Resource*, *Request* and *Response*, along with the fields to represent those types. An instance of each type can be created to represent actual resources, responses and requests. Interactive forms to create, view, edit and delete the content types have been created along with fine-grained

---

<sup>2</sup> Available from <http://drupal.org/project/rdfx>

<sup>3</sup> Modules available from <http://drupal.org/project/modules>

multi-user permissions for doing the same. Mappings between each content type and the ontology have also been created, so that *Resource*, *Request* and *Response* content are presented as open data.

Screen-shots of the prototype application are shown in figure 2. The main navigation menu (a) accesses features such as searching for resources (b), creating requests (c) and viewing responses (d) from the network of sites.

In order to ensure meaningful URLs, we have developed an additional module for Drupal, entitled Pathauto for RDF <sup>4</sup> By default Drupal URLs are not meaningful; so the URL for a blog post might be `http://example.com/?q=node/1` or `http://example.com/node/1`. However, with the built-in URL rewriting feature and the Path Auto module, content URLs are made readable, such as `http://example.com/blog/example-title`. The effect of this is URLs that can be used as URIs for the items stored. However, as the Path Auto and RDF modules are not fully compatible, the RDF document describing an item is available at the meaningless URL, `http://example.com/node/1.rdf` even if the web page is a Path Auto one. The Path Auto for RDF module automatically creates an alias for the RDF document at the same location as the URI e.g. `http://example.com/blog/example-title.rdf`.

#### 4.1 A Decentralised Network

So far we have described a stand-alone online system for managing resources. In principle, the system so-far is equivalent to the kinds of resource management systems used by emergency management organisations. Next, we describe how this is used as the basis for a decentralised network of co-operative sites using a novel linked-data based mechanism. Two ontologies have been created: Protocol for Request-Response (PRR) and an Ontology for Real Resources (ORR)<sup>5</sup>, to respectively describe requests and responses and the physical resources that can be transferred in a network of sites. The Ontology for Real Resources (ORR) is used to encode information about resources as RDF triples that are interchanged between organisations to describe resources in a common conceptual framework.

In order to establish a network, two steps are required. Firstly, two or more sites must exist, each openly publishing resources, requests and responses. The sites can include displays for viewing openly published resources, requests and responses from the network. In Drupal, the displays can be created using Views created using the the SPARQL Views module, to show internally held data and data from external SPARQL endpoints. Secondly, any two sites can form cooperative network by subscribing to one another; adding each other as SPARQL data source. Subscription requires site-users to add external sites. This allows sites to:

- View and search for resources under the remit of external sites

---

<sup>4</sup> Available under the terms of the GPL from <http://drupal.org/sandbox/shah/1871288>.

<sup>5</sup> These ontologies form part of a wider effort in the Disaster 2.0 project to survey and fill the gaps in ontologies for disaster and emergency response cf. [10]

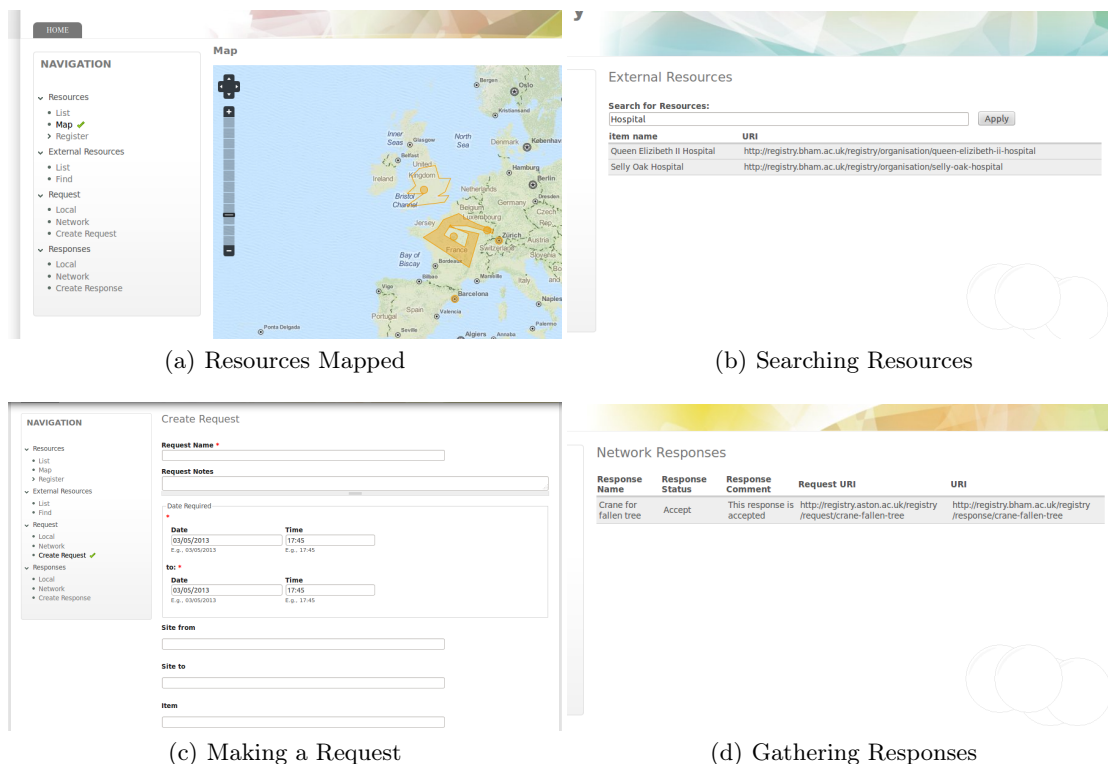


Fig. 2. Prototype user interface screen-shots

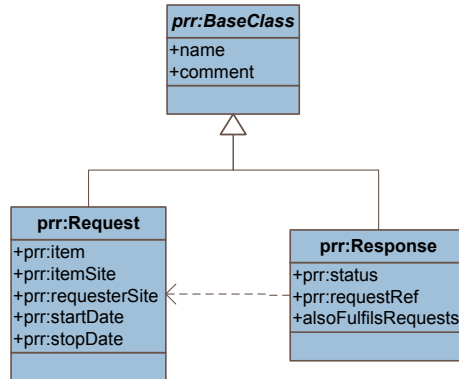
- Publish requests for resources from an external site
- View the responses for requests made to an external site
- View requests made for the resources under the remit of the local site
- Publish responses to requests for local-site resources

In effect creating an interoperable network of distributed independent sites for resource management in an emergency. The following sections describe the linked data based protocol and the ontology for describing resources.

## 4.2 PRR: Protocol for Request Response

The Protocol for Request Response ontology consists of three classes and ten properties, shown in figure 3. The two main classes are `prr:Request` and `prr:Response`, which extend from the abstract `prr:BaseClass` for the common properties provided. Sites publish instances of either `Request` or `Response` for a given `Resource` to enact the protocol.

Given two sites (e.g. *bham*, *aston*) that hold a register of resources and the sites subscribe to each other, the protocol is as follows. The *bham* site makes a



**Fig. 3.** Class diagram: Protocol for Request Response

request for a widget from *aston*, by creating and publishing an instance of the `ppr:Request` class. The `ppr:Request` slots are filled as shown in Table 1.

Slot	Value
Request Name	Widget Request
Comment	With our thanks
Item	http://aston/widget
Item Site	http://aston/
Requesting Site	http://bham/
Start Date	[desired date & time]
Stop Date	[desired date & time]

**Table 1.** A request by *bham* for a widget from *aston* with URI `http://bham/request/widget-request`, this request is encoded as triples and published on the *bham* site.

The request is published as open data and accessible via the *bham* SPARQL endpoint. As the *aston* site subscribes to the *bham* endpoint, the request appears on the *aston* site. The *aston* site can either accept or reject the request for the widget by creating an instance of the `ppr:Response` class that references the original `ppr:Request` from *bham*. The *bham* site will see the response as it subscribes to the *aston* site.

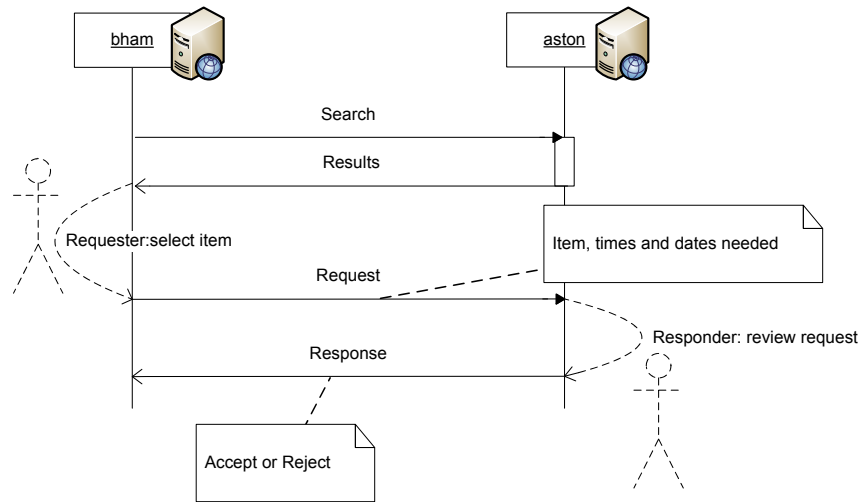
The request is published as linked data and accessible via the *aston* SPARQL endpoint. As the *bham* site subscribes to the *aston* endpoint and the response contains the URI of the original request, the *bham* site can determine if the original request was accepted. In the case the request is not accepted, the comment slot can help the requester re-submit an acceptable request. The protocol steps in this scenario are outlined in the sequence diagram shown in Figure 4.

<sup>6</sup> Where there are similar requests made by other sites for the same need, they are linked in this slot.



Slot	Value
Response Name	Widget Response
Comment	Not a problem
Request URI	http://aston/request/widget-request
Status	“accepted”
Also Fulfils Requests	Requests fulfilled by this response <sup>6</sup>

**Table 2.** A response by *aston* to the request for a widget from *bham* with URI `http://aston/response/widget-request` this request is encoded as triples and published on the *aston* site.



**Fig. 4.** Sequence diagram: overview of Protocol for Request Response

## 5 Evaluation via Case Study

We evaluated the system by means of a conceptual evaluation based on emergency management organisations. Various emergency services, local and national government are the responders. Requests can include items such as sand bags, sand, water pumps, and pontoons.

In this type of emergency a great amount of co-operation and co-ordination is needed between the emergency responders. Each emergency service needs to respond to the needs of the affected population with available resources. An important issue with such response is interoperability. As noted above, traditionally each organisation is only aware of their own resources, and the only way they are aware of resources of what exists elsewhere is through informal networks, telephone calls etc. Each emergency service typically uses a separate system for managing resources, as the needs and response of each service naturally differ. Software systems which are designed as stand alone and without interoperability cause problems when responders require resources from outside

organisations. Responders must find and contact the relevant person to request resources, and making manual requests for resources can take valuable time away from response.

With the proposed system in place, each emergency service runs a separate instance, for example, the fire service controls the site 'FireServices' and the local council controls the 'LocalCouncil' site. The sites publish data openly via the mechanism set out in this paper but are normally used for managing resources within the organisations. The two sites subscribe to each other and form a network of independent but cooperative sites for resource management. This fulfils several of the evaluation criteria. Firstly the systems allow for 'open information exchange', publishing requests with the principles of linked open data. The sites are obviously decentralised and interoperable using a direct open data protocol based on RDF and SPARQL. The independent systems are also expected to be in regular use for day-to-day resource management.

In an emergency, such as a flood scenario, FireServices responders can use the system to request resources from the Council site and vice-versa. For example, if residents are being evacuated, the LocalCouncil site can send a request to the FireServices site for rafts by publishing the request openly. This eliminates the need for the council to find and contact the correct person, instead the user has 'Information Focus' on getting the resources needed for the evacuation task, using familiar tools. 'Crisis Memory' is also built into the systems, as previous requests and responses are available from each organisation in the network.

The proposed systems cannot make crises more predictable, so it does not directly address the issues 'exceptions as norms' or 'crisis scope' mentioned above in Section 2. However, because of the flexibility of the system, any site can request from any other, allowing for 'Unpredicted Coordination'. So for example, a local builders' merchant can set up a site to join the network and respond to requests to provide sand bags for the affected population. Finally, roles are transferable in the systems at both macro and micro levels; responders can also become requesters and individual users can be granted more privileges to respond to and make requests. Because the system uses open data, it is possible for all organisations involved to have the most up-to-date information about resources available, for greatest "information validity".

This initial demonstrative evaluation has shown how the proposed system fulfils the aim set out in this paper and the criteria of [3]. In order to perform a complete evaluation, an emergency practice exercise is planned with the Inter-municipality Civil Protection Center<sup>7</sup> in the Province of Pisa and Livorno, Italy.

## 6 Related Work

Several works from the literature are related to the current proposed system for cooperative resource management in emergency. The protocol used for communication between sites forms a Semantic Web Service [11], however in this work,

---

<sup>7</sup> <http://valdicecina.salaoperativaprociv.org/>

the focus is on physical resource exchange and so a minimal protocol is created as a very-lightweight semantic web service. Furthermore, the current tool uses geo-location information and semantic web in a similar vein to stSPARQL [12] and the closely related geoSPARQL [13]. In the current work, simple latitude and longitude pairs are used to represent resource locations. Further development is required for the integration of geoSPARQL or stSPARQL like features.

More closely related to the prototype presented here are Ushahidi [5] and Sahana Eden [6] tools. As previously noted, these tools are centralised with single installations and not-interoperable with other installations of the same tool. [14] describe a system for presenting existing data for emergencies as Semantic Web Services, this system is also centralised and external. Sahana Eden publishes data openly via a RESTful API [15], however no applications have been found to exploit this data. In the current prototype, cooperation and interoperability between sites using the same software is a key feature.

Another area of related work are the ontologies for managing disasters: Management of a Crisis Vocabulary (MOAC) [16] and Humanitarian eXchange Language (HXL) [17]. MOAC terms originated from the Ushahidi platform and use during the Haiti crisis, HXL is inspired by MOAC. The aim for both ontologies is to present and exchange information about disasters, attempting to cover the entire disaster. In contrast, the ontologies presented here focuses on a specific aspect of disaster- resource exchange. Furthermore, MOAC and HXL are based on observations of crises and have not been integrated into a tool that publishes data, whereas the ORR and PRR ontologies are developed for inter-site cooperation and interoperability the prototype tool and to achieve the goals from [3].

## 7 Conclusion

This paper has presented an ontology and protocol for resource management in emergency response. A network of sites is created by mutual subscription to SPARQL endpoints, where available resources are published openly. Any site in the network can also use the open data based protocol to publish requests for particular resources. The site to which the resource belongs can respond via the protocol. The proposal consists of a distributed network of cooperative but independent organisations that can publish and exchange resources, without centralised control. The proposed protocol is also designed to interoperate with existing an legacy systems.

## Acknowledgement

This research has been supported by the Disaster 2.0 project (<http://www.disaster20.eu>) funded by the ‘Prevention, preparedness and consequence management of terrorism and other security-related risks’ programme, European Commission – Directorate-Home Affairs (HOME/2010/CIPS/AG/002).

## References

1. United-Nations: United Nations International Strategy for Disaster Risk Reduction, UNISDR Terminology on Disaster Risk Reduction (2009)
2. Fiedrich, F., Gehbauer, F., Rickers, U.: Optimized resource allocation for emergency response after earthquake disasters. *Safety Science* **35**(1) (2000) 41–57
3. Turoff, M., Chumer, M., Van de Walle, B., Yao, X.: The design of a dynamic emergency response management information system (DERMIS). *JITTA* **5**(4) (2004) 1–35
4. Currión, P., Silva, C.d., Van de Walle, B.: Open source software for disaster management. *Communications of the ACM* **50**(3) (2007) 61–65
5. Okolloh, O.: Ushahidi, or 'testimony': Web 2.0 tools for crowdsourcing crisis information. *Participatory Learning and Action* **59**(1) (2009) 65–70
6. Careem, M., De Silva, C., De Silva, R., Raschid, L., Weerawarana, S.: Sahana: Overview of a disaster management system. In: *Information and Automation, 2006. ICIA 2006. International Conference on, IEEE* (2006) 361–366
7. Shadbolt, N., Berners-Lee, T., Hall, W.: The Semantic Web Revisited. *IEEE Intelligent Systems* **21**(3) (2006) 96–101
8. Berners-Lee, T.: *Linked Data - Design Issues*. Technical report (2006)
9. Heath, T., Bizer, C.: *Linked Data: Evolving the Web into a Global Data Space*. Morgan Claypool (2011)
10. Liu, S., Shaw, D., Brewster, C.: *Ontologies for Crisis Management: A Review of State of the Art in Ontology Design and Usability*. In: *Proceedings of the Information Systems for Crisis Response and Management conference (ISCRAM 2013 12-15 May, 2013)*. (2013)
11. McIlraith, S.A., Son, T.C., Zeng, H.: Semantic web services. *Intelligent Systems, IEEE* **16**(2) (2001) 46–53
12. Koubarakis, M., Kyzirakos, K.: Modeling and querying metadata in the semantic sensor web: The model stRDF and the query language stSPARQL. *The semantic web: research and applications* (2010) 425–439
13. Open Geospatial Consortium: *OGC GeoSPARQL-A geographic query language for RDF data*. OGC (2011)
14. Tanasescu, V., Gugliotta, A., Domingue, J., Davies, R., Gutiérrez-Villarías, L., Rowlatt, M., Richardson, M., Stinčić, S.: A semantic web services GIS based emergency management application. *The Semantic Web-ISWC 2006* (2006) 959–966
15. Richardson, L., Ruby, S.: *RESTful web services*. O'Reilly Media, Incorporated (2007)
16. Ortmann, J., Limbu, M., Wang, D., Kauppinen, T.: Crowdsourcing linked open data for disaster management. In: *Proceedings of the Terra Cognita Workshop on Foundations, Technologies and Applications of the Geospatial Web in conjunction with the ISWC*. (2011) 11–22
17. Keßler, C., Hendrix, C., Limbu, M.: *Humanitarian eXchange Language (HXL) Situation and Response Standard* (2012)