IST-Africa

IST-Africa 2006 Conference Proceedings
Paul Cunningham and Miriam Cunningham (Eds)
IIMC International Information Management Corporation, 2006
ISBN: 1-905824-01-7

openAdap.Net: A Technical Perspective

Alessandro E.P. VILLA[1], Javier IGLESIAS[1,2], Solange GHERNAOUTI-HELIE[3]
[1]*INSERM U318 Université Joseph Fourier, Grenoble Cedex 8, 38043, France*
Tel: +33 4 7676 5625, Fax: + 33 4 7676 5619, Email: avilla@neuroheuristic.org
[2]*INFORGE Université de Lausanne, HEC, Lausanne, 1015, Switzerland*
Tel: +41 21 692 3587, Fax: + 41 21 692 3585, Email:Javier.Iglesias@unil.ch
[3]*INFORGE Université de Lausanne, HEC, Lausanne, 1015, Switzerland*
Tel: +41 21 692 3421, Fax: + 41 21 692 3405, Email: sgh@unil.ch

Abstract: This paper presents an Internet infrastructure, called **openAdap.Net**, that makes possible the dissemination of simulation resources, and their exposure to application and evaluation across domains in ways that might be unanticipated. Simulation processing tools issued from physical sciences could permeate to study problems as different as the dynamics of societal interactions, linguistic analyses, crops forecast, traffic congestions, and life sciences. **openAdap.Net** is aimed to break current boundaries to resource sharing, and hence supports transdisciplinarity. End-users are provided with the ability to browse and apply shared resources, and dynamically compose and integrate existing resources to leverage new research insights.

Keywords: Resource sharing, Internet protocols, data processing, open source, software engineering, Internet computing.

1. Introduction

The information circulating in the Cyberspace generates a historically unprecedented richness for sharing knowledge and methods of data processing. The software published at Internet in one instance becomes available for the whole society worldwide. However, most of the latest processing and simulation tools developed and provided by research laboratories and businesses are based on specific software and computer platforms.

The current situation generates a lag until these contributions become known within the same field of competence and restrict severely their availability, in particular for cross-fertilization application in other fields. This delay may provoke the "re-invention" of methods for data processing and, more generally, the "re-discovery" of the same knowledge. At the societal scale this delay inhibits the development of added-value activities originating from sharing the knowledge and generates repeated unnecessary expenses and even, erroneous applications. The societal benefits in terms of business developments, market diversification and creation of employment generated by the fast circulation and ease of access to the results of the Human Genome Project illustrate the interest in developing an open and dynamic adaptive network for resource sharing with emergent properties within the Information Society.

openAdap.Net stands in the area of complexity and aims to provide a distributed environment where simulation tools of all kinds (applications, data, knowledge, …) can be accessed transparently via the Internet. At this time, three architectures are used in this field: Grid, Web-services, and Peer-to-Peer (P2P). The peculiar strongholds of these architectures are briefly described here.

Grid: *"Each user has a large dataset to manipulate with one application distributed on a set of computers."* The problem addressed by the Grid is to distribute the processing of large data sets. The required application is copied to distinct computers over a network with each treating a portion of the data. The results are recomposed from all the partial results at the end of the treatment. End-users have control on both data and applications, but little information on remote execution hosts.

Web Services: *"Many users exploit the same services permanently provided through a centralized authority."* Web Services provide a secured and controlled access to applications, usually to collaborators in the same company or institution. The goal is to provide distributed access to the same references, databases or applications. The system is articulated around a repository where the service interfaces are published in a self-contained form. The architecture is rather static, as services seldom change and are expected to be permanently available. End-users have no control over the applications and little information on remote execution hosts.

P2P: *"Many users exchanging pieces of data in an unsupervised way."* P2P (peer-to-peer) systems address the data-sharing problem. Copies of the applications installed on end-users' computers keep open connections from one computer to peers', forwarding queries and results back and forth until a match is found somewhere on the network. The architecture is open and competing implementations coexist on a self-organized network. End-users have control over their data and information on the peer hosts. It is interesting to note that end-users tolerate incredibly poor service quality and that this architecture raises social as well as technical issues.

openAdap.Net falls somewhere between these three architectures exploiting several of those interesting aspects, but with the intention to address a two-way problem: To provide to a community of users in the same domain a means to interchange their resources in an open, dynamic and secured way and to provide to a community of users the access to the exploitation of information processing solutions contributed by users belonging to other communities. End-users have control over their data, but do not need to manage the resources, nor do they have complete information on remote execution hosts. Collaboration within the openAdap.Net network allows the dynamic integration of these resources, possibly yielding new or previously unexpected composite resources. This can be summarized as follows: *"Many users interchanging resources (data, applications, knowledge, …) dynamically provided by interconnected domain-oriented brokers."*

The above description is synthesized in Table 1.

Table 1: Comparison of openAdap.Net with the other approaches over Internet

	Data treatment distribution	Hardware resource allocation	Hidden execution hosts	Applications sharing	Published application interfaced	Data sharing	Highly dynamic system	Transparent connections between users and resources
Grid	X	X	X					
Web Services			X	X	X	X		
P2P						X	X	X
openAdap.Net	X	X	X	X	X	X	X	X

2. Objectives

The project will develop further an openAdap.Net demonstrator (located at `http://www.openAdap.net/`). The present demonstrator is an ongoing development that possesses the full end-to-end range of features required to deliver usability to

contributors, end-users and potential exploiters, thus allowing proof of concept validation and subsequent development. To achieve the final goal of openAdap.Net a number of objectives must be met:

- **User friendliness:**
  Human-machine interfaces will be designed to speed the acceptance of openAdap.Net. Convenient simulation tools will be provided for broker setup and management, resource sharer setup and management, convenient web-portals and openAdap.Net-aware applications within case studies. Innovative solutions to control, use and manage the emergent properties of distributed systems such as openAdap.Net will be investigated. Viewing openAdap.Net as a virtual operating system will help finding comprehensible solutions that users with poor knowledge in computer science can easily manage by transposing what they have previously learned. Particular emphasis will be placed on data visualization.

- **Interoperability:**
  openAdap.Net is heavily based on XML technologies to define an interoperable and loosely coupled architecture capable of dealing with dynamic runtime appearance and disappearance of resources from the network. Resources will be described in such a way that all components will be able to manipulate them, whichever form they take. No specific network hardware or protocol will be required to ensure the system will be able to operate with forthcoming infrastructures (wireless, high-speed, IPv6, …). A new URI name scheme `oan://` will be proposed to the WWW Consortium. A complete API will be proposed to the Open Source community so that they can contribute to the open architecture.

- **Security and Privacy:**
  Focussing on security within open environments means to define the targets of threats. Without doing a risk analysis survey, the main targets of security threats are: end-user, network access point, network and all the infrastructures connected to the network as servers and information systems (see Figure 1) [1].

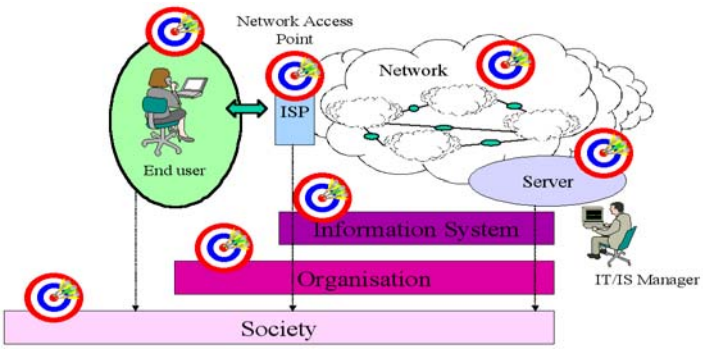

Figure 1: Security within open environments.

The negative impacts of security threats will affect not only systems or individuals actors but in a chained way organisations and society [1]. For OpenAdap.net the security is addressed transversally by defining rules on how to handle data at any point in the system, providing the required profiles for identification and authentication of the entities, and offering strongly encrypted connections when communicating through the Internet. Standards and models of the field will be applied from the start of the project. The use of unconditional secure encryption scheme (such as quantum cryptography) is also possible [2].

- **Adaptability:**
  Brokers handle the complex task of routing queries and results between brokers, resource sharers and end-user openAdap.Net entry points. Bio-inspired mechanisms will be applied to improve the decision making process, using neural network learning to identify common bottlenecks for example and adapt the broker behaviour to the circumstances, while genetic mechanisms could be used to evolve broker communities and let emerge new properties.

- **Open Implementation:**
  A reference implementation will be proposed to test the validity of the concept. Case studies will be set up to apply openAdap.Net to specific domains of competence. A generic web portal will be developed that will be reused for the other case studies. Already existing applications will be modified to become openAdap.Net-aware. The infrastructure will be open so that other contributors will be given the opportunity to build their own communities, web portals and openAdap.Net-aware applications. Collateral repositories will be set up for domain-specific ontologies and benchmark data for each case study. Reconfigurable hardware will be built to optionally access the openAdap.Net network.

- **Dissemination:**
  Conference demonstrations and presentations will be performed. Journal articles will be published on the outcomes of the project. A dedicated website will be available for knowledge transfer.

3. Methodology

The openAdap.Net project is an advanced engineering system, which is:

- **Complex**
  Resources and the composition of such can be of arbitrary complexity. The interactions between components in the openAdap.Net network, required for the processing of an end-user request can be of arbitrary complexity.

- **Distributed / networked**
  Resource hosts and brokers need only be connected to the Internet, and registered with openAdap.Net.

- **Heterogeneous**
  Resource sharers, brokers, and end-user platforms may all employ different software platforms.

- **Autonomous**
  Brokers will act autonomously on behalf of end-users in deciding how to process user requests.

- **Adaptive**
  The broker network will self-adapt in order to optimise the performance of the openAdap.Net system.

- **Semi-intelligent**
  Employ symbolic reasoning and adaptation.

The project implementation is designed around 8 Work Packages. The goal of WP1 is to build the ontologies describing the openAdap.Net project itself and its components (brokers, resources, etc.) and a specific ontology for the purpose of naming conventions.

Our original approach is to treat openAdap.Net as a virtual distributed operating system. The ideas collected from this paradigm will enrich the Application Program Interface (API) with high-level functionalities in a consistent way. Nevertheless openAdap.Net goes beyond the traditional operating systems in that the complete system configuration is highly volatile. WP1 2 deals with all communication related tasks (messaging API, specification of a new name scheme `oan://`, and implementation of the communication libraries). The work packages combine pure research (e.g., adaptation mechanisms for a network of brokers, WP3), technological design (e.g., processor design, WP5), and iterative prototypical development of generic openAdap.Net software components (resource sharer tools WP4; brokers WP3; databases WP3; simulation of the dynamics through graph theories WP7) and end-user community specific artefacts (portals WP6 and ontologies WP1). The management of the project and the contact with SME and Third Parties is the task of WP8.

4. Technology Description

4.1. Project components

***Resources**:* In principle we place no limits on what might constitute a resource – it might for example be an analysis tool, a data repository, benchmark files, or an ontology. Resources must be hosted at a resource sharer (and registered with openAdap.Net there-from), although in response to system adaptation, such resources may be partially duplicated at various sites on the openAdap.Net network. Resources are provided to the system by a *resource contributor*.

***Resource sharers**:* The resource sharer is the resource host. In the case of a resource providing say an analysis tool, then the host providing execution; in the case of a data repository, the host acts simply as a server of the data. Resources available at a particular resource sharer are (previously) registered at one of more brokers, who can then request that the resource sharer serves up the resource, possibly with input and output data being passed between the broker and the resource sharer. Resources sharers can also provide meta-information to enquiring brokers, such as (i) information about load capability, both for the resource sharer as a whole, and for specific resources; this information can be used by the resource sharer to inhibit the allocation of further resource requests, and (ii) meta-information about specific resources such as data (input and output) formats and ontological description.

Resource sharers also control information to/from brokers such as (i) reservation (by a broker) of a particular time-slot for resource provision, or (ii) failure recovery information. The main difference between openAdap.Net and Grid Computing in resource management is in the fact that with openAdap.Net the association of an analysis with a resource is done statically, by registration. We can assume in the openAdap.Net case that the application could be considered trusted more easily than an application downloaded from a remote location (Grid case). In order for an application to be registered, it would have to pass a "stress test" on its invocation parameters, to ensure that no easily apparent security holes are detected. This stress test could be available on the Web, similarly to the way one could check HTML pages for W3C conformance.

Security issues have to be considered more from a user perspective. Resources could have different levels of trust with a given user according to user-specific criteria, in terms of data handling (encryption, suppression in case of failure, …), resource identification (location, type, …), or execution control (stop analysis, …).

***Adaptive broker**:* The network of brokers forms the first layer of the openAdap.Net middleware, being responsible for the (dynamic) processing of end-user requests. Brokers will also have the ability to self adapt.

***openAdap.Net-aware applications**:* An openAdap.Net-aware application might consist of an end-user interface to one or more openAdap.Net resources, provided for a single, or group of *end-users*. The project presents examples of visualisation tools that might be deployed in this way, for example allowing the user to view multi-dimensional output data from an analysis tool.

***End-user community web portals**:* A web portal provides more flexibility in that a greater number of resources are gathered together as are appropriate for a particular end-user community, and described in terms of the terminology of that community. Examples of such portals will be developed in four diverse, but important, application areas. Portals may also provide the end-user with the ability to dynamically compose and integrate resources, via a visual programming interface. The use of portals by end-users will be supported by evolvable portal-specific ontologies.

***openAdap.Net Managers and System Administrator**:* In addition to openAdap.Net contributors and end-users, openAdap.Net Managers and System Administrators will require access to system information. This will be provided by individual resource sharers and brokers, and a visualisation tool will be provided for the analysis and monitoring of the information interchange within the network components. Figure 2 shows the overview of project components and their interactions.

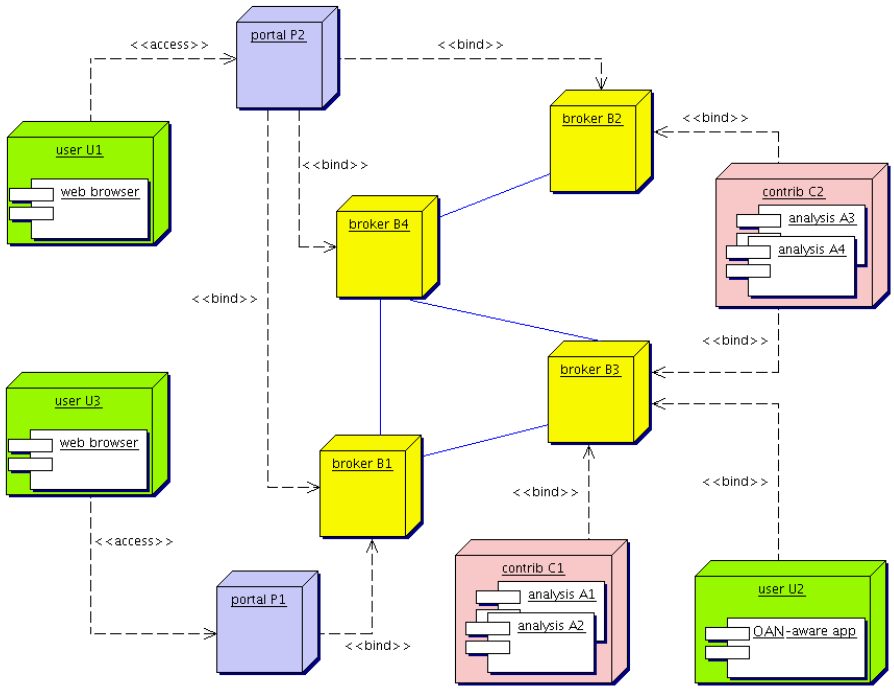

Figure 2: Overview of openAdap.Net components and their interactions.

4.2. Contribution to the standards

A Uniform Resource Locator (URL) [3] is a compact string representation of the location for a resource that is available via the Internet. URLs are designated by including a "<scheme>:" and a "<scheme-specific-path>". Many URL schemes are already defined, like "`ftp:`", "`http:`", and "`mailto:`" in the Internet Engineering Task Force (IETF) [4] tree.

As part of the documentation process, it is foreseen to submit an open specification for a URL scheme name "oan:" [5,6] as an Informational (non-standard) Request for Comments (RFC) [7] to the IETF, requesting it to be registered in the IETF tree of URL scheme names. This will open up the key feature of resource localization on openAdap.Net networks and give the opportunity to anyone to provide his own implementation or low-level contribution to openAdap.Net core.

In order to register "oan:" as a new URL "<scheme>" name, the syntax and semantics of the "<scheme-specific-path>" has to be properly defined, described and commented. IETF also insists on the scheme name to be implemented and extensively tested before submitting it. If enough interest is driven by "oan:" in the Internet community, the Informational RFC might be promoted to a Standards Track RFC, leading to the establishment of an Internet Standard [8].

5. Developments

openAdap.Net is an intelligent network infrastructure supporting the use of shared resources, such as data, knowledge, tools and expertise aimed at studying and exploiting the emergent properties from dynamic interacting agents.

Individual users of openAdap.Net can be classified as either contributors (of shared resources), or end-users of such. In addition to individual use, openAdap.Net is open to exploitation by networked organisations and alliances, by providing an infrastructure with which they can share and integrate resources, and develop new business solutions and opportunities. In the context of commercial exploitation, security is a vital issue and this will be an intrinsic part of the openAdap.Net specification and design.

From the viewpoint of a contributor, openAdap.Net makes possible the dissemination of resources, and their exposure to application and evaluation by the wider user community. The support for wider evaluation of simulation tools and frameworks is of course particularly important in the research arena, as is the ability of other researchers to reproduce computational results. openAdap.Net also allows the dissemination of resources across domains in ways that might be unanticipated by the original contributor. For example a solution developed by experts in the field of medical statistics could permeate to the field of botany, and become a key tool in the analysis of tree distributions. openAdap.Net thus breaks current boundaries to resource sharing, and hence supports transdisciplinarity.

End-users are provided with the ability to browse, and apply shared resources. openAdap.Net also supports the user in dynamically composing and integrating existing resources, broadening the data interoperability obtained with XML to all types of resources. Such ability may be employed to leverage new research insights, or to provide new opportunities for business exploiters of openAdap.Net. In addition to these immediate benefits to the end-user, there is also a benefit to the user-community in general, in that the amount of time spent redeveloping the same tools and knowledge is an important limitation to the acquisition of new ones.

openAdap.Net is an Open Source project designed as an open architecture, meaning that anyone is invited to contribute their own enhancements to the system, via a set of libraries and tools that will be provided by the consortium. Such an initiative is aimed at increasing the impact of the project with all the ideas we haven't yet foreseen that competent enthusiasts will imagine and realize for their specific domains. We will also contribute to the open standards with a URI name scheme (oan://) that will be used to locate resources over the network, and possibly with a dedicated protocol.

Internally, the network of openAdap.Net brokers will be responsible for decomposing and routing end-user tasks to appropriate resource sharers for execution. The manner in which this decomposition and routing takes place will be determined dynamically taking

into account loading, availability and other issues. The required negotiation between brokers (and resource sharers) may be compared to agent interaction. The openAdap.Net network is also able to self-adapt via learning processes derived from neural and artificial life learning. Such learning might result in new broker-broker connections, reassessments of the value of such connections, specialisation or generalisation or broker behaviour, etc. This is definitely a plus to the existing architectures for distributed environment like grid and web services that points out the income expected from the project to the networked computing field.

6. Results

The openAdap.Net project is not directly aimed at the production of new methods of analysis. In a first step, contributors will be invited to aggregate their already existing tools for the community to use them. In a second step, openAdap.Net will become an environment where new techniques and methodologies will gain access to a wide range of users, possibly over the direct community boundaries to the adjacent domains.

One major impact on the scientific community could be obtained by attracting well known scientific editors to encourage the scientific authors to provide source data to the community repositories, and to share their methods through a openAdap.Net web portal, hence promoting reproducibility of results, and direct peer review of the published article methodologies and results by the readers.

One particular exploitation (possibly via SMEs), is the development of web portals tailored to the needs of end-user communities centred on specific domains. We would expect this to give rise to synergies within and between both new and existing communities. The identification of such communities will necessitate further developments such as community specific ontology and benchmark file repositories.

openAdap.Net being an open source project, contributions are expected from yet unknown developers and interested professionals after one year from project start, when public events and public publishing of openAdap.Net outcome will make openAdap.Net established in the Cyberspace. This will be encouraged by setting up the full range of community development tools and techniques like CVS – Concurrent Versions System – servers, bug trackers, publicly accessible. There is enough experience by now on the impact of such coordinated and distributed development schemes on the quality of the resulting products to embrace it.

We foresee four potential partnerships to be integrated into the running project:

- The adaptive and behavioural models for the broker implementation represent major innovations of the openAdap.Net project. We believe that several SMEs and research laboratories will be interested in developing novel interdisciplinary solutions associated to the psychological and technological aspects of evolvable simulation tools, the psychological environment of remote user support and the formal aspects of artificial processing in resource-sharing.
- Interfacing openAdap.Net with other distributed systems like grids and clusters might require adapters to be developed, such as openAdap.Net to PBS – Portable Batch System – interfaces.
- Interactions with Publishers and Editors could have an incredible impact on the way scientific dissemination, including visualization, is performed, mainly by facilitating peer review in refereed publications and for results comparison and validation for peer scientific readers. Such interactions could be envisioned once the openAdap.Net system will be fully available to the scientific community and should be encouraged as an extension of the case studies.

- Privacy and Security issues are rapidly evolving in the Cyberspace. We expect potential new partners in this field which provides extensive possibilities of participation to SME.

7. Business Benefits

The ability to tackle a scientific problem from a new perspective relies on both the past experience and new skills adopted by an individual. The openAdap.Net project is based on the collaboration between information scientists, electronic engineers, computer scientists and neuroscientists having diverse scientific interests and very specialized techniques. We feel that such a transdisciplinary approach is a necessary way for the achievement of real advances in producing impacts in the Information Society Technologies. The strengthening of collaborative links as a result of undertaking this work will allow the promotion of an inter-disciplinary research outlook, essential for a successful scientific career in the rapidly developing middleware. It is envisaged that exchange visits between Post-Doctoral workers from the different partner countries will form an integral part of the work plan – such exchanges will increase the scientific and cultural outlook of these young scientists and generate an increased learning experience.

The openAdap.Net architecture is applicable to real-life case studies where we will build communities in the domains of competence of the partners to assess the validity of the concepts and the usability of the implementation. Each of the case studies will take the aspect of tailored web portals that end-user will use to access the resources shared by contributors. We expect those portals to be integrated in the set of tools that each of those domains typically share, and to be more and more undertaken by the communities at the end of the funding by this project. Users will be invited to contribute resources, articles, benchmark data, and to build the domain ontology with the initial help of openAdap.Net promoters.

One of the by-products of this research will be the opening up of information technology systems to a wider audience through its applications. Through this contribution, users of any background can play a more active role associated with interactive systems, emphasising positively their position in a *spectrum of ability*. There is a real sense that in defining new modalities in interactive systems, everyone feels in some sense 'disabled', as anyone who attempts to play an unfamiliar musical instrument will realise. Such new active role that everybody may play promotes a social policy of *inclusion*. The adaptation and reconfiguration offered by systems implemented using openAdap.Net will provide an enabling aspect to all users by latching onto their individual strengths and weaknesses; a process that could become fully automated.

Citizens of less favoured countries will have access to all shared openAdap.Net resources with a basic Internet connection, thus benefiting from the knowledge transfer and available assets, and contributing back to the community with their own approaches and resources. The outcome of the expected cross-fertilization is unpredictable. Side effects are expected on the quality and harmonization of the resource documentation, as an important effort is dedicated to the elaboration of tools to enhance them.

The ability to evaluate and to find the best available solution to a given problem may have significant impact on areas such as economics, physics, environmental sciences, meteorology, and health. By adopting the best available method of analysing a set of data, it is possible to test alternative solutions and choose the best from its overall performance. In the construction industry, which is playing an important role measured by investments in all economical systems, even small enhancements will make resources available to other purposes. Reduction of the energy consumption during the whole lifecycle of a building will decrease pollution and save money. By setting up an environment where people are used to working with state of the art technical solutions, companies throughout the World will be forced to increase their competitiveness work in a completeness way.

Pushing existing paradigms like neuronal network inspired learning rules for the adaptable information processing or the operating system paradigm for the overall communication layout, and the lessons learned for 10 years on the self, dynamically, openly organized content on the web over the TCP/IP adaptive network are key aspects of the openAdap.Net philosophy and architecture for resource sharing.

8. Conclusions

This paper has presented the main features of openAdap.Net, which is an intelligent network infrastructure supporting the use of shared resources, such as data, knowledge, tools, expertise, etc. aimed at providing the most advanced tools for data analysis and simulation to a broad audience over Internet with limited resources to information processing. Individual users can be classified as either contributors of shared resources, or end-users of such. openAdap.Net will be open to exploitation by networked organisations and alliances taking into account the vital issue of Internet security and privacy.

The openAdap.Net infrastructure makes possible the dissemination of simulation resources, and their exposure to application and evaluation across domains in ways that might be unanticipated. Simulation processing tool issued from physical sciences could permeate to study problems as different as the dynamics of societal interactions, linguistic analyses, crops forecast, traffic congestions, and life sciences. openAdap.Net is aimed to break current boundaries to resource sharing, and hence supports transdisciplinarity. End-users are provided with the ability to browse and apply shared resources, and dynamically compose and integrate existing resources to leverage new research insights.

The openAdap.Net is an Open Source project increasing its impact by contributing to the open standards with an URI name scheme (`oan://`) that will be used to locate resources over the Internet, and possibly with a dedicated protocol that will be specified and implemented also in a scalable hardware demonstrator. The modelling schemes of openAdap.Net will be implemented using specific hardware substrate with autonomous growth, adaptation and fault-tolerance capabilities. This hardware will be organised around a regular array of programmable basic elements and with a specialised subsystem able to manage complex I/O structures aimed at simplify complex, computationally intensive and off-line compilation process. openAdap.Net brokers will be responsible for dynamically decomposing and routing end-user tasks to appropriate resource sharers for execution. The negotiation between brokers (and resource sharers) is inspired by the way how the brain processes information. The openAdap.Net network will be able to self-adapt via learning processes that could give rise to modifiable broker-broker connections, specialisation or generalisation of broker behaviour, etc. The nonlinear dynamics that will emerge from our approach makes openAdap.Net closer to the complexity of a living organism.

Acknowledgments

This work is partially funded by the European research project SECOQC [`http://www.secoqc.net/`].

References

[1] S. Ghernaouti-Hélie. Guide de cybersecurité pour les pays en développement, ITU publication, 2006.
[2] S. Ghernaouti Hélie, M.A. Sfaxi, G. Ribordy, O. Gay. Using Quantum Key Distribution within IPSEC to secure MAN communications. MAN 2005 conference, 2005.
[3] RFC 2396 - Uniform Resource Identifiers (URI): Generic Syntax.
[4] http://www.ietf.org/
[5] RFC 2717 - Registration Procedures for URL Scheme Names.
[6] RFC 2718 - Guidelines for new URL Schemes.
[7] RFC 2223 - Instructions to RFC Authors.
[8] RFC 2026 - The Internet Standards Process -- Revision 3.

[9] Konstantin Läufer. A hike through post-EJB J2EE web application architecture, Computing in Science and Engineering, 7(5):80–8, 2005.
[10] G. Eisenhauser, F.E. Bustamante, K. Schwan. Publish-subscribe for high-performance computing, IEEE Internet computing, 10(1):40–7, 2006.
[11] X. Zhang, S. Chen, R. Sandhu. Enhancing Data Authenticity and Integrity in P2P systems, IEEE Internet computing, 9(6):42–9, 2005.