Network Domain Federation - Infrastructure for Federated Testbeds

Sebastian Wahle[1], Anastasius Gavras[2], Fabricio Gouveia[1], Halid Hrasnica[2], Thomas Magedanz[1]

[1]Fraunhofer FOKUS, Berlin, Germany; [2]Eurescom, Heidelberg, Germany;

E-mail: [1]{sebastian.wahle|fabricio.gouveia|thomas.magedanz}@fokus.fraunhofer.de, [2]{gavras|hrasnica}@eurescom.de

Abstract: **This document describes a generic approach on how to federate networking domains. Network federation is a model for the establishment of a large-scale and diverse infrastructure for communication technologies, services and applications. A federation of network environments can generally be seen as an interconnection of two or more independent network domains for the creation of a richer environment and for the increased multilateral benefits of the users of the individual domains. This paper discusses testbed federation as a concrete example for network domain federation and outlines a technical architecture framework as the result from the IST FP6 Panlab Specific Support Action.**

Keywords: Network Domain Federation, Federated Testbeds, IST, FP6, Panlab, Teagle

1 INTRODUCTION

The main driver for the idea of *network domain federation* and federated testbeds is that dedicated testing and network infrastructures tend to be very expensive. Professional network equipment, such as carrier grade high speed network elements, is highly resource consuming both in initial investments and maintenance costs. To illustrate the possible impact of the network domain federation principle, consider as an example that for every dollar spent on a new server, companies tend to spend up to 50 cents on electricity and cooling. Having this in mind and considering the latest discussions on Green IT, it can be said that the IT (Information Technology) industry is facing a shift from over-provisioning to virtualization and *Infrastructure as a Service* (IaaS) concepts.

IaaS, being one of the emerging buzzwords after *Service Oriented Architectures* (SOA) and *Software as a Service* (SaaS), means that customers no longer purchase infrastructure (both servers/network equipment and software), but acquire such resources as an outsourced service. This is in line with the main idea of this paper which investigates the requirements for interconnected network domains and providing composite infrastructures that are spanning all technology layers from network connectivity to service architecture and that can be used by both industry and academia on demand in a combinational manner.

In a network domain federation, the domains are usually geographically dispersed and owned by different organizations. They would however be considered as being part of a single entity (virtual environment), in so far as they are operated in a common management framework under a common management authority, provided by a network domain federation. The federations are dynamic and evolve over time based on the requirements of the users. However, the operation and management of federated environments over multiple networks and administrative domains is difficult and requires specific mechanisms. There, particular areas of interest are mechanisms and tools to describe, store, locate and orchestrate services and infrastructure components as well as automatically provide virtual composite network environments across multiple administrative domains, in order to provision various services established by federating different network domains.

The provisioning and management of highly heterogeneous infrastructures, such as the proposed federated networking landscape, is challenging from a technical point of view. So far, the management of distributed environments has been approached by unifying as much as possible of the underlying infrastructure. However, the idea for the central coordination entity defined here is to impose minimum overhead to the owners of resources and their customers, thus a fine balance must be found between efficiency and fine grained management. A representative example for possible application of network domain federation, investigated in this paper, is the establishment of a testbed federation.

The *IST FP6 Panlab SSA* [1] defined a framework for the interconnection of independent testbeds and laboratories so as to facilitate the establishment of a *Pan-European Laboratory*. The concept of such a Pan-European Laboratory is based on the federation of distributed testbeds that are interconnected and provide access to required platforms, networks and services for a broad range of activities such as demonstrate, trial and evaluate new technologies, system solutions and telecom service concepts.

Published in NEM Summit 2008 proceedings, 13-15 October 2008, http://www.nem-summit.eu/

Corresponding author: Sebastian Wahle, Fraunhofer FOKUS, Kaiserin-Augusta-Allee 31, 10589 Berlin, sebastian.wahle@fokus.fraunhofer.de

This paper focuses on the technical infrastructure necessary for implementing and maintaining the Panlab concepts, serving as a case study for the more generic approach of network domain federation.

This paper is organized as follows: The second section outlines the concept of testbed federation and describes Panlab roles, structure and development phases. The third section addresses interconnection issues within the federation. Section four describes the central federation tool as part of the overall architecture and lists technologies for implementing some of the required functions. Finally, the main conclusions are drawn.

2 CONCEPT FOR TESTBED FEDERATION

The approach of network domain federation is rather generic and defines a concept for two or more cooperating network domains to achieve a common goal. This is useful for many activities such as outsourcing, prototyping, proof-of-concept realization and testing. However, in the following we will describe the use case of *testbed federation* and a framework for building an on-demand infrastructure of federated testbeds rather than elaborating on the generic approach for network domain federation.

2.1 Panlab Idea – Federation of Testbeds

As already touched in the introduction, the Panlab [1] SSA defined concepts for achieving federation of distributed testbeds and to interconnect them in order to provide access to different platforms, networks and services. Furthermore, objective of Panlab is to pave the way for a clear view of where in Europe is the best place to test a specific technology, system, service or application. The implementation of a physical infrastructure will be enabled, aiming at establishing integration, testing and validation services for product prototypes in Europe and beyond.

The overall vision is to establish the grounds for a future operational and long-term self-sustainable Pan-European laboratory including a business model that would advise on the continuation of the action as an independent entity. There, customers shall be enabled to design, configure and provision a desired infrastructure on demand, making use of automated processes and tools offered by the Panlab organization.

2.2 Panlab Structure

The proposed Panlab testbed federation consists of multiple testbeds offered by *Panlab Partners* that are interconnected through a Virtual Private Network (VPN), the *Panlab Office* and supporting tools such as the *Teagle* (search and orchestration engine).

This is shown in Figure 1, where the following Panlab roles have been identified:

- The Panlab Partner represents an entity responsible for providing the testbed facilities used by Panlab Customers.
- The Panlab Office is an entity in charge of coordinating and supporting the Panlab organization as well as managing technical, business and marketing activities.
- The Panlab Customer uses the resources provided by the Panlab Partners and the Panlab Office, to test new technologies, components, system solutions and telecom service concepts.

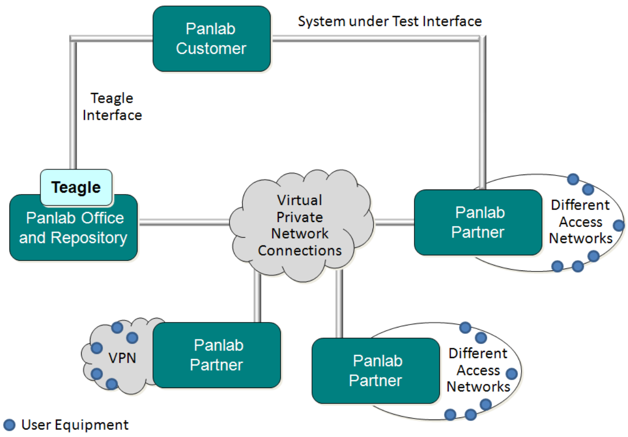

Figure 1: Basic Panlab Architecture

The technical infrastructure model involves the use of the following components to achieve connectivity and uniformity among Panlab Partners, Panlab Customers and the Panlab Office [2]:

- Connectivity component: is to be considered as VPN-based tools to enable the interconnection between Panlab Partners. These tools may be presented as resources themselves, hiding the technical complexities and allowing Panlab Partners to have multiple overlay networks for testbeds on "as required" basis.
- Teagle: is the federation search and orchestration engine. Teagle will provide a smart interface for Panlab Customers making available necessary information and means for serving their needs. That is a web-based service to fulfil Panlab Customers needs related to Panlab offerings. This interface should provide access to a database containing information on testbeds and exploit a repository of testbed functionality descriptions. This interface will also be able to combine different testing components into a functional entity making it possible to meet various testing requirements that a single testbed may not be able to satisfy.
- *Panlab Repository*: is a component to maintain the description of the Panlab Partner testbeds, detailing the testing services and caching resources offered by each Panlab Partner to the Panlab community. It shall also be able to maintain test results and configurations for repeating previous experiments.

2.3 Phased Panlab Deployment

The technical infrastructure offered by Panlab is not static. It is a rather dynamic structure that will evolve, including new platforms, services and systems offered by previously connected or new laboratories.

The infrastructure will be deployed in stages, evolving from a simple to more advanced configurations. The Panlab evolution path will go from a manual phase (where testbed descriptions are manually entered by testbed representatives) to a more automated solution allowing both Panlab Partners and Panlab Customers to set-up and configure the required testbeds on demand.

3 ENSURING CONNECTIVITY IN THE FEDERATION

In order to offer composite infrastructure environments consisting of elements from different administrative domains, there is the need to fix some common interconnection mechanisms that will be obeyed from all participating domains. At the same time, one of the main Panlab objectives is to impose minimum requirements on the individual domains. This is a dilemma as all domains collectively have to agree on some common mechanisms while at the same time Panlab concepts try to abstract as much as possible from any limiting processes in order to maximize flexibility. The solution proposed here is to make use of so called *Modular Gateways.* Such gateways, as shown in Figure 2, reside at the border of a testbed.

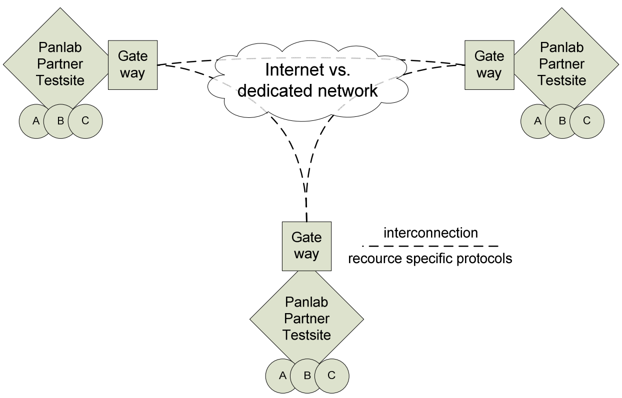

Figure 2: Panlab Connectivity

The gateways are acting as signalling converging points translating federation level signalling to any resource specific communication. Also, the gateways are responsible for setting up dynamic VPN links. The central federation control unit, in this case the Teagle tool, communicates VPN set-up requests towards the gateways, which are responsible for providing the requested links and assuring connectivity to the components needed for the demanded testbed infrastructure.

VPN technology provides the means for setting up secure overlay networks on top of possibly unsecure network links. Currently, it has not been decided whether Panlab should rely on dedicated network links or make use of the open Internet to interconnect Panlab Partners and the Panlab Office.

Another aspect regarding connectivity is of course *addressing*. All the participating testbeds represent individually functional environments that have their own customers, ongoing project activities, etc. Therefore, all components inside the Panlab Partner testbeds dispose of their own addresses while they are merely used on the Panlab Partner testbed level. However, as soon as a component is booked from a Panlab Customer in order to run an experiment on the Panlab federation level, the addresses of the involved components presumably need to change to be addressable on the federation level. A common addressing scheme for the Panlab federation is needed and components need to be addressable accordingly. The gateways are responsible of communicating this addressing change to the testbed components upon demand.

4 TEAGLE – ENABLING TOOL FOR FEDERATED TESTBEDS

It is clear that the challenging task of building a testbed federation spanning multiple countries, network boundaries and administrative domains, requires certain control mechanisms and entities. While centralized approaches have lately been challenged by distributed peer-to-peer approaches, certain functionalities are very difficult to provide, following a distributed approach. Among those functionalities is for example authentication. It can be said that especially industry organizations are somewhat reluctant to opening cooperate infrastructures for the proposed idea of testbed federation. Usually, such infrastructures reside behind well-configured firewalls and access is limited and controlled. As Panlab especially aims at integrating industry environments, a trustworthy relationship needs to be build in order to convince as many industrial partners as possible to participate. This leads to an important design decision. Panlab relies on a centralized approach where much functionality is provided by centrally administered tools and the central Panlab Office. Nevertheless, the mentioned centralized approach does not exclude the possibility that parts of the infrastructure could be placed and maintained in a decentralised manner if appropriate.

Besides technical issues, centralized Panlab Office services include legal and operational aspects (contract templates, etc.). We believe that trust can be build best with many partners if there is a central *Federation Control* unit offered by a central *Business Entity* (see Figure 3) that, in case something goes wrong, can also be held liable for what was contractually agreed beforehand.

The following sections introduce the federation architecture and describe some of the most promising technologies to achieve the objective of building a testbed federation.

4.1 Architecture

Figure 3 shows a general architecture for network domain federation where the federation consists of interconnected network domains that offer certain services and components depicted by the circles A, B, C and D. The central *Business Entity* provides the *Federation Control* unit and service composition tools where the services and components offered by the domains can be orchestrated on demand. In a top-down approach the interconnected domains can be configured and managed, while the domains publish their services and capabilities bottom-up.

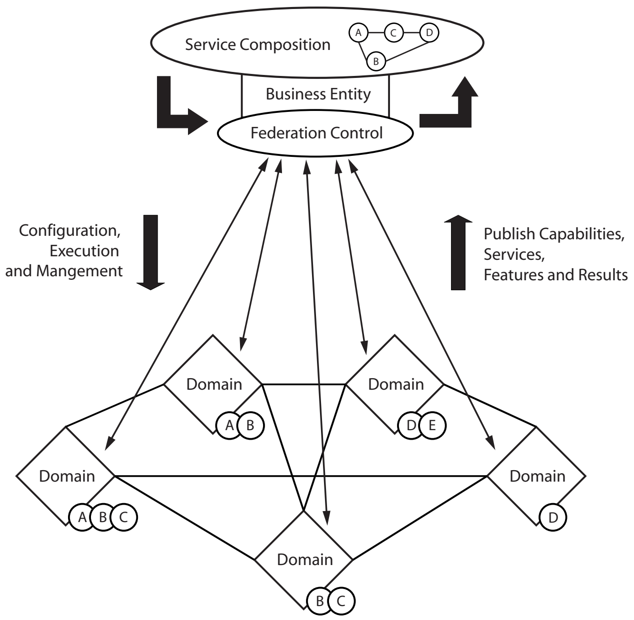

Figure 3: Federation Architecture

This is now mapped to the use case of testbed federation as shown in Figure 4. The network domains are represented by testbeds that are provided by the Panlab Partners. The central federation control unit is represented by Teagle and will be described in more detail throughout the following section.

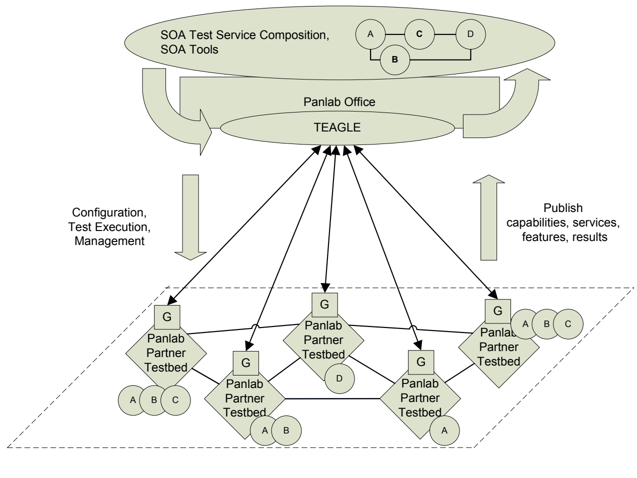

Figure 4: Testbed Federation

The central business Entity is represented by the Panlab Office that provides technical, legal and organisational support to Panlab Customers and Panlab Partners.

4.2 Technologies

The high-level architecture, presented above, requires a number of concepts and technologies. This section provides an overview of existing technologies that have been identified to provide a good starting point for the realization of the proposed testbed federation system. Figure 5 shows the functional building blocks that need to be implemented. For most of these blocks we will outline existing technologies and standards that mostly derive from the *Web Services* [3] and *Semantic Web* [4] area. As shown in Figure 5, the main parts are a *repository*, an *ontology* definition, an *orchestration* component and a service *broker*.

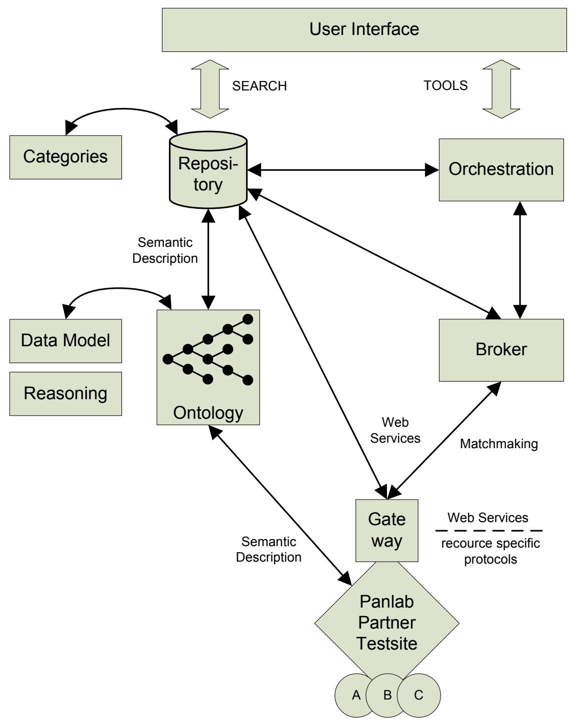

Figure 5: Teagle Functional Building Blocks

There is the need for a uniform way to describe all services and components offered by the Panlab Partner testbeds (this includes semantic descriptions). The descriptions need to be stored in the repository. The orchestration and broker components make use of the descriptions in order to enable the composition of services and components and perform the execution of composite services and provisioning of the desired environment.

Web Services provide a well-known approach for machine-to-machine communication across the boundaries of networks and administrative domains. This makes them highly suitable for the federation level communication. Usually, Web Service communication means transporting *Simple Object Access Protocol* (SOAP) [5] messages via the *Hypertext Transfer Protocol* (HTTP) [6]. SOAP is a lightweight protocol for exchanging data between computer systems. This usually results in transporting *Extensible Markup Language* (XML) [7] data.

As said before, a *semantic description* of the services and components offered is required.

4.2.1 Federation Ontology

Ontologies provide a means for formally describing and defining a domain. Usually, entities within a domain and the relationships between entities are formally represented. Furthermore, ontologies can be used to reason about certain properties of the domain. Currently, the *Web Ontology Language* (OWL) [8] is widely used to produce and publish ontologies. For the proposed testbed federation, ontologies shall be used to define the testbed offerings and provide the necessary semantics for meaningful descriptions. As building comprehensive ontologies from scratch is a difficult and time consuming exercise, it is foreseen to reuse existing work for describing the telecommunication domain. For example, the Tele Management Forum (TMF) [9] *Shared Information Data* (SID) model [10] and the IST FP6 *SIMS* [11] *ontology* [12] already provide a solid knowledge base and might be transferred into a suitable ontology.

The SID model was born from the need for achieving interoperability on sharing and reusing data from different components in the communication industry. It is a common language developed by the TMF that is used for defining common data and provides a knowledge base that is used to describe the behaviour and structure of business entities as well as their collaboration and interactions.

4.2.2 Semantic Descriptions

Once a suitable ontology has been build to represent the testbed offering, the *terminology and semantics* provided by this ontology can be used to produce descriptions of the offerings that are to be stored in the repository. The idea is that the Modular Gateways that reside at the border of each testbed provide a number of Web Services that can be used to set-up and configure testbed offerings or that provide a testbed-specific service. The *Web Service Description Language* (WSDL) [13] can be used to describe such Web Services. However, WSDL can merely describe syntactic elements of a Web Service, that means how a client can access a specific service, while semantic information cannot be exposed using plain WSDL. Therefore, *Web Service Semantics* (WSDL-S) [14] and *Semantic Annotations for WSDL* (SAWSDL) [15] have been specified, that allow for semantic annotations (using terms and semantic concepts defined by an ontology) of existing WSDL files. Another approach is to link the WSDL files for concrete services to the ontology using *Semantic Markup for Web Services* (OWL-S) [16].

4.2.3 Repository

Once the services and components offered by the testbeds have been (semantically) described, the descriptions need to be stored in a repository for later retrieval. An available technology for such a repository is a *Universal Description, Discovery and Integration* (UDDI) [17] registry. UDDI registries allow for exposing information about a business (or other) entity and its technical interfaces or Application Programming Interfaces (API). Semantic information can be linked to UDDIs as well as categorization information for components and services. The *United Nations Standard Products and Services Code* (UNSPSC) provides an open standard for accurate classification that could be used to classify and categorize Panlab offerings.

4.2.4 User Interface

Another functional building block shown in Figure 5 is the *User Interface*. This is a critical element for the entire architecture and the Panlab vision in general, as it must be intuitive and easy to handle. The User Interface shall allow the lookup of available testing technologies, components and services. Furthermore, it shall allow the expression of a request for a desired infrastructure and possibly a graphical tool to design the infrastructure based on the Panlab Partners offerings. This is where most of the before mentioned functionality is required in a single place.

To enable search functionality, the User Interface must be connected to the repository. Among the possible search options are *free text search* and *guided search*. The later enables the user to choose from displayed options. Once a top level option has been chosen, the user will be prompted for more specific (lower level) options and details, increasing the granularity of the result.

4.2.5 Service Orchestration

Orchestration shall enable the user to design a desired infrastructure as a composition of Panlab Partner offerings. Therefore, the semantically described and registered infrastructure components and services need to be represented graphically and provided as drag and drop objects on a virtual sketch board. From a technology perspective orchestration could for example be realized using structured OASIS *Web Services Business Process Execution Language* (WSBPEL) [19] sequences.

4.2.6 Service Brokering

The *Broker* box also shown in Figure 5 is needed to invoke the service request sequence (that is received from the orchestration component) and to perform availability checks if the desired infrastructure can be provisioned. The Broker is a component that might evolve from a rather simple matchmaker component to a more advanced policy enforcement and request delegation component. As the Panlab implementation will follow a phased approach, the broker can evolve during time and offer additional functionality in later stages. Currently, the area of policy evaluation, policy enforcement and policy management in *Next Generation Networks* (NGN) is still subject to research and standardization [20].

5 CONCLUSIONS

Acknowledging some simple facts, such as that a single testbed cannot provide every possible testing environment or every possible testing resource, or that testing resources such as high guaranteed bandwidth network links or dedicated testing equipment are very expensive, lead to the concept of network domain federation and the specific proposal to federate existing testbeds. By doing so, the scattered available resources become available through a single logical entry point, which increases visibility and potential utilization of expensive resources.

For realizing the proposed Panlab federation, a supporting technical infrastructure must be developed and deployed that takes into account basic technical requirements; among others:

- How to interconnect remote testing resources
- How to describe and locate testing resources
- How to enable access in a uniform way

The necessity to interconnect different, potentially remote, testing resources for the provisioning of a specific testing environment is a challenging task. The diversity of potential testing resources that could be federated and made available via the proposed infrastructure is very high, allowing the comparison that its complexity exhibits similarities with the complexity of controlling converged networks in the context of next generation networks and services (NGN).

The large variety and different character of available testing resources that must be made visible, accessible and comparable through a single logical interrogation point leads to a set of technical requirements, which relates to the way how the testing resources are described, stored, discovered, brokered and provisioned. To satisfy this set of requirements and as a *first conclusion*, this paper proposes the implementation of a new tool, called Teagle, which unifies the representation of testing resources for the purpose of later automated processing. Resources under consideration in this area include testing components, environment configurations, test suits, prototyping services and testing results.

The *second conclusion* is that research work has to be carried out, in order to identify and develop a suitable representation by re-using and extending existing state-of-the-art description, orchestration and brokering techniques.

The *third conclusion* is that the Panlab technical infrastructure must grant controlled access at different levels to different resources, allowing for resource allocation, scheduling, access rights management, as well as access to information related to management of the federation for the purpose of fulfilling service level agreements and generally to assure quality of offerings. In order to satisfy these requirements in a controlled way, research and development work must be carried out to build on existing resource abstraction techniques.

Overall, the Panlab technical infrastructure must impose minimum overhead to the owners of testing resources and their customers, thus a fine balance must be found between efficiency and fine grained management, so that "market dynamics" could be induced in the European testbed landscape. The Panlab vision is that the market dynamics will work towards state-of-the-art high quality testbed offerings in the federation.

Acknowledgments

Parts of this work have been developed in the IST FP6 Panlab SSA which has been funded by the European Commission's Sixth Framework Programme. The authors of this paper would like to thank the Panlab members for the good collaboration.

References

[1] The Panlab project website, http://www.panlab.net/
[2] Fabricio Gouveia and Sebastian Wahle (Eds.), Approach to technical infrastructure, IST FP6 Panlab Deliverable, March 2008, http://www.panlab.net/publications/public-deliverables.html
[3] The W3C Web Services Activity, http://www.w3.org/2002/ws/
[4] The W3C Semantic Web Activity, http://www.w3.org/2001/sw/
[5] The Simple Object Access Protocol specification, http://www.w3.org/TR/soap/
[6] The Hypertext Transfer Protocol specification, http://www.ietf.org/rfc/rfc2616.txt
[7] The W3C Extensible Markup Language publications, http://www.w3.org/XML/Core/#Publications
[8] The W3C Web Ontology Language specification, http://www.w3.org/TR/owl-features/
[9] The Tele Management Forum website, http://www.tmforum.org
[10] The Tele Management Forum Shared Information Data model http://www.tmforum.org/browse.aspx?catID=2008
[11] The IST FP6 Semantic Interfaces for Mobile Services (SIMS) project website, http://www.ist-sims.org/
[12] The SIMS ontologies, http://www.tele.pw.edu.pl/~sims-onto/
[13] The W3C Web Services Description Language Version 2.0 specification, http://www.w3.org/TR/wsdl20/
[14] The W3C Web Service Semantics (WSDL-S) member submission, http://www.w3.org/Submission/WSDL-S/
[15] The W3C Semantic Annotations for WSDL (SAWSDL) website, http://www.w3.org/2002/ws/sawsdl/
[16] The W3C Semantic Markup for Web Services (OWL-S) specification, http://www.w3.org/Submission/OWL-S/
[17] OSASIS Universal Description, Discovery and Integration specifications, http://www.oasis-open.org/committees/uddi-spec/doc/tcspecs.htm
[18] The United Nations Standard Products and Services Code, http://www.unspsc.org/
[19] The OASIS Web Services Business Process Execution Language http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsbpel
[20] The OMA Policy Evaluation, Enforcement and Management specification, http://member.openmobilealliance.org/ftp/public_documents/arch/Permanent_documents/