Network Domain Federation - An Architectural View on How to Federate Testbeds

Sebastian Wahle, Prof. Dr. Thomas Magedanz
Fraunhofer FOKUS, Berlin, Germany
sebastian.wahle@fokus.fraunhofer.de

This article describes a generic approach on how to federate networking domains. Network domain 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. The article discusses testbed federation as a concrete example for network domain federation and outlines a technical architecture framework.

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 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 server/network equipment and software), but acquire such resources as an outsourced service. This is in line with the main idea of this article which outlines some 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. Federations are dynamic and evolve over time based on the requirements of the users. 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. 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.

Federation Connectivity

The approach of network domain federation is rather generic and defines a concept where two or more cooperating network domains engage in order to achieve a common goal. This is useful for many activities such as outsourcing, prototyping, proof-of-concept realization and testing. 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 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 we 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 *Gateways*, as shown in the figure below, that reside at the border of a domain. The gateways are acting as signalling converging points translating federation level signalling to any resource specific communication. Also, the gateways are responsible for establishing dynamic VPN links.

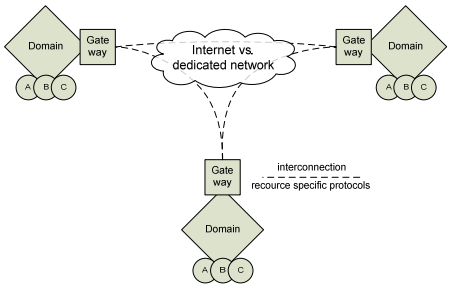

Federation Connectivity

A central federation control unit communicates VPN set-up requests towards the gateways, which are responsible for providing the requested links and assuring connectivity to the domain resources. VPN technology provides the means for setting up secure overlay networks over possibly unsecure network links. This is possible on top of dedicated network links or the open Internet.

Enabling Tools for Federation

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. Therefore, a trustworthy relationship needs to be build to convince as many industrial partners as possible to participate. This leads to an important design decision. The architecture relies on a centralized approach where much functionality is provided by centrally administered entities and tools.

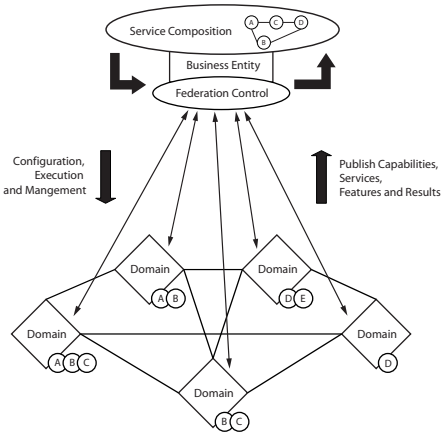

Federation Architecture

Trust can be build best with many partners if there is a central *Federation Control* unit offered by a central *Business Entity* (see figure above) that, in case something goes wrong, can also be held liable for what was contractually agreed before. In the general architecture for network domain federation shown above, 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.

This high-level architecture requires a number of concepts and technologies.

Published as Position Statement for the FIREweek, 10 – 12 September 2008, Paris, France,
http://www.ict-fireworks.eu/events/fireweek-in-september.html

The figure below shows the functional building blocks that need to be considered. The main parts are a *Repository*, an *Ontology* definition, an *Orchestration* component and a *Service Broker*.

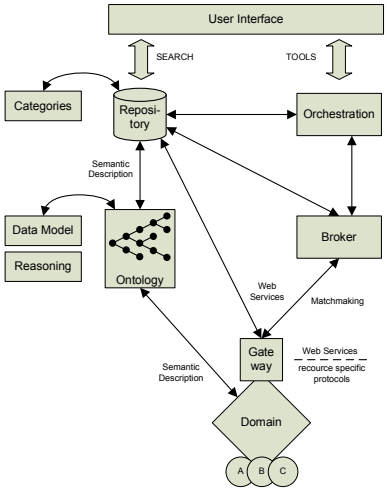

Functional Building Blocks

There is the need for a uniform way to describe all services and components offered by the testbeds (this includes semantic descriptions). The descriptions need to be stored in the repository. The orchestration and broker components make use of the descriptions to enable the composition and invocation of services. *Web Services* [1] 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) [2] messages via the Hypertext Transfer Protocol (HTTP) [3]. SOAP is a lightweight protocol for exchanging data between computer systems. This usually results in transporting Extensible Markup Language (XML) [4] data.

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) [5] 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) [6] Shared Information Data (SID) model [7] and the IST FP6 SIMS [8] ontology [9] already provide a solid knowledge base and might be transferred to a suitable ontology.

Semantic Descriptions

Once a suitable ontology has been build to represent the testbed offerings, 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 *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) [10] 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) [11] and Semantic Annotations for WSDL (SAWSDL) [12] 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) [13].

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) [14] registry. UDDI registries allow for exposing information about a business (or other) entity and its technical interfaces or Application Programming Interfaces (APIs). 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) [15] provides an open standard for accurate classification that could be used to classify and categorize the domain offerings.

User Interface

Another functional building block shown in the figure above is the *User Interface*. It shall allow the lookup of domain resources such as 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 available domain resources. 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.

Service Orchestration

Orchestration shall enable the user to design a desired infrastructure as a composition of available domain resources. Therefore, the semantically described and registered infrastructure components and services shall 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) [16] sequences.

Service Brokering

The *Service Broker* 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. The area of policy evaluation, policy enforcement and policy management in *Next Generation Networks* (NGN) is still subject to research and standardization [17].

Acknowledging 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, endorses the network domain federation concept and the specific proposal to federate existing testbeds. By doing so scattered domain resources become available through a single logical entry point, which increases visibility and potential utilization of expensive resources. This article highlighted the main issues to be addressed and motivates the practical implementation of the ideas described.

References

[1] W3C Web Services Activity, http://www.w3.org/2002/ws/
[2] SOAP specification, http://www.w3.org/TR/soap/
[3] HTTP specification, http://www.ietf.org/rfc/rfc2616.txt
[4] W3C XML publications, http://www.w3.org/XML/Core/#Publications
[5] W3C OWL specification, http://www.w3.org/TR/owl-features/
[6] Tele Management Forum (TMF) website, http://www.tmforum.org
[7] TMF SID model http://www.tmforum.org/browse.aspx?catID=2008
[8] IST FP6 SIMS project website, http://www.ist-sims.org/
[9] The SIMS Ontologies, http://www.tele.pw.edu.pl/~sims-onto/
[10] W3C WSDL Version 2.0 specification, http://www.w3.org/TR/wsdl20/
[11] W3C WSDL-S submission, http://www.w3.org/Submission/WSDL-S/
[12] W3C SAWSDL website, http://www.w3.org/2002/ws/sawsdl/
[13] W3C OWL-S specification, http://www.w3.org/Submission/OWL-S/
[14] OSASIS UDDI specifications, http://www.oasis-open.org/committees/uddi-spec/doc/tcspecs.htm
[15] UNSPSC website, http://www.unspsc.org/
[16] OASIS WSBPEL specification, http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsbpel
[17] The OMA PEEM specification, http://member.openmobilealliance.org/ftp/public_documents/arch/Permanent_documents/

Published as Position Statement for the FIREweek, 10 – 12 September 2008, Paris, France, http://www.ict-fireworks.eu/events/fireweek-in-september.html