Available online at www.sciencedirect.com
SCIENCE DIRECT
Future Generation Computer Systems 22 (2006) 862–867
FGCS FUTURE GENERATION COMPUTER SYSTEMS
www.elsevier.com/locate/fgcs

Using RDF to describe networks

Jeroen J. van der Ham[a,*], Freek Dijkstra[a], Franco Travostino[b], Hubertus M.A. Andree[a], Cees T.A.M. de Laat[a]

[a] *Advanced Internet Research Group, Universiteit van Amsterdam, Kruislaan 403, 1098 SJ Amsterdam, Netherlands*
[b] *Nortel Labs, Boston, MA, USA*

Available online 3 May 2006

Abstract

Conventions such as iGrid 2005 and SuperComputing show that there is increasing demand for more service options on networks. For such networks, large teams of experts are needed to configure and manage them. In order to make the full potential of hybrid networks available to the ordinary user, the complexity must be reduced.

This paper presents the idea of the Network Description Language (NDL), which builds on Semantic Web techniques to create a distributed Topology Knowledge Base (TKB). The TKB can provide a collection of reachability graphs, showing connectivity rules among physical and/or virtual entities.

Latching onto the Semantic Web provides network management with a new breed of tools—bots, compilers, browsers, both commercial off-the-shelf (COTS) and open source. The approach appears to be applicable to the Global Lambda Integrated Facility (GLIF) as well as other experimental communities.

Keywords: Network descriptions; Semantic Web; Resource Description Framework; Hybrid networks

1. Introduction

A hybrid network is a challenging engineering feat: they offer end users both traditional IP services and new optical services in the form of dedicated lightpaths. The main purpose of hybrid networks is to navigate efficiently and effectively through the service options offered by Layers 3, 2, and 1. In such a network, islands that peer at Layer 3 will advertise reachability information along with service options that map to DiffServ Code Points [1]. Other islands that operate at Layer 2 will need to indicate, for instance, which one of the 10 Gigabit Ethernet framing techniques they use (e.g., LAN/WAN PHY). Also, islands that operate at Layer 1 add a host of TDM and/or WDM options.

Enter Grids. Grids make a strong case for specialisation around a plurality of network services. Regular network traffic is typically many-to-many and is optimally handled within a packet-routed network. Conversely, the super-sized traffic is often few-to-few and is best handled via agile Layer 1 services, as implemented in SURFnet [2], the research network of The Netherlands. The latter services yield a circuit-switched network experience with uncontested bandwidth and strong non-interference properties. A desirable Grid Network is therefore a network wherein packet-routed segments and Layer 1 and 2 “cut-through” segments across different clouds can be assembled end-to-end by a discerning user or, better yet, a software agent on his behalf. The degree of user control is a sharp departure from inter-carriers or carrier-of-carriers agreements oblivious to specific traffic demands. In a Grid Network, a user (or software in his behalf) is seen playing an active role in the negotiation of service stipulations.

This paper presents the Network Description Language (NDL), which builds upon the Resource Description Format (RDF) [3] and its linking capabilities to produce a distributed Topology Knowledge Base (TKB) [4]. The advantage of the TKB is that it provides easily accessible knowledge that the management and control planes can build upon. This is especially true for the dynamic provisioning planes, such as UCLP [5] or DRAC [6].

* Corresponding author.
E-mail address: vdham@science.uva.nl (J.J. van der Ham).

0167-739X/$ - see front matter © 2006 Jeroen van der Ham. Published by Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2006.03.022

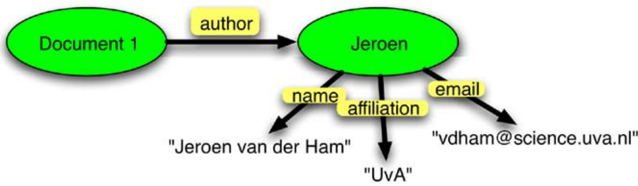

Fig. 1. A simple RDF graph.

This paper continues with a short introduction to the technology that NDL builds upon, RDF and the Semantic Web, in Section 2. Then, in Section 3, the Network Description Language is introduced. An overview of future research is given in Section 4 and the conclusion is in Section 5.

2. Introduction to RDF and the Semantic Web

The World Wide Web allows us to publish and share documents with other people in the world. However, because it has become so popular and so widespread, it has almost become the victim of its own success. The large scale and the abundant availability of data make it very hard to find what we want. Search-machines, such as Google or Yahoo, have indexed the data and provide searching services. However, computers still have no common sense, so the search capabilities of the search machines are rather limited. Consider, for example, the following two sentences:

- *A* is connected to *B*.
- There is a connection between *A* and *B*.

These two sentences may have a different meaning. There is no way for a computer to understand that these two lines have the same meaning. This is where the Semantic Web can help computers (and people). The following is an excerpt from the activity statement of the Semantic Web initiative [7]:

> *The goal of the Semantic Web initiative is to create a universal medium for the exchange of data where data can be shared and processed by automated tools as well as by people. For the Web to scale, tomorrow's programs must be able to share and process data even when these programs have been designed totally independently.*

In 2000, the Semantic Web initiative was started by the World Wide Web Consortium (W3C). Since then, the W3C has been working on several specifications to publish and share (meta)data, including the Resource Description Framework (RDF) [3] and SPARQL [8]. In the following section, we provide a brief introduction to RDF. Examples of RDF and SPARQL are given later on in Section 3.

2.1. Resource Description Framework

The Resource Description Framework is a method for representing information about resources on the Web. It provides a common framework for expressing metadata so that it can be exchanged between applications without loss of meaning.

Information in RDF is expressed using triplets, with the following elements:

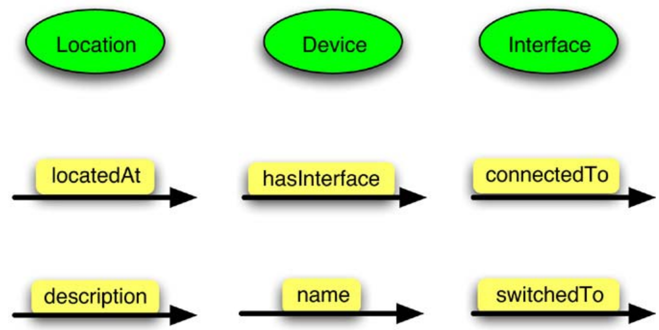

Fig. 2. Overview of the classes and properties of the Network Description Language.

Subject. The resource being described
Property. The property that the statement describes
Object. The value of the property according to this statement.

A set of triplets is called a graph. Using the property that an object can also be the subject of another triplet, complex graphs can be created. An example of such a graph is shown in Fig. 1. This graph shows that `Document1` is written by `Jeroen`. It also provides additional information about him, such as his name, affiliation and email address.

The graph in Fig. 1 has the same problem as the two lines shown before. We have provided an abstract way of defining relations, but use plain English as labels for identifying these relations. Consider the `author` relationship; we could have expressed this as `creator` without loss of meaning to human readers. RDF solves this terminology problem by using URIs.[1] Related terms are usually defined using the same URI-prefix, taking the form of XML namespaces. See, for example, the Dublin Core Metadata Initiative [10].

There are several ways of expressing RDF graphs. One is the graphical form as in Fig. 1. The most common textual form is RDF/XML [11], where the graph is encoded in an XML format. Another, more compact format is *Notation3* [12]. Throughout this paper we use the RDF/XML notation, which allows us to use tools for XML as well as RDF. Examples and explanation of the RDF/XML syntax are given in the next section.

An application of RDF is the Friend of a Friend (FOAF) [13] project. FOAF is similar to what we are trying to achieve with describing networks. That is, a distributed network of descriptions, managed by the responsible stake-holders, that together form a complex, navigable network. Tools are available for creating FOAF descriptions [14], or to browse the network created by the links in the FOAF descriptions [15].

3. Network Description Language

The distributed descriptions and emergent network properties of the FOAF project are what we need for network descriptions. This is why we set out to create a simple ontology in RDF to describe physical networks, the Network Description Language. An overview of the classes and properties of the Network Description Language is given in Fig. 2.

[1] URI stands for Uniform Resource Identifier. A URL is one kind of URI. See also [9].