Future Generation Computer Systems 42 (2015) 55–63

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Resource discovery and allocation for federated virtualized infrastructures

C. Pittaras [a], C. Papagianni [b,*], A. Leivadeas [b], P. Grosso [a], J. van der Ham [a], S. Papavassiliou [b]

[a] System and Network Engineering research group, University of Amsterdam, The Netherlands
[b] NETwork Management and Optimal DEsign Laboratory, National Technical University of Athens, Greece

HIGHLIGHTS

- Networking innovations over virtualized infrastructures—NOVI federation architecture.
- Resource discovery and allocation over multi-domain virtualized infrastructures.
- Prototype of a distributed, semantic-based resource discovery and mapping framework.
- A semantic aware virtual network embedding algorithm.
- The approach is evaluated via simulation against non-semantic aware solutions.

ARTICLE INFO

Article history:
Received 1 August 2013
Received in revised form
6 November 2013
Accepted 18 January 2014
Available online 1 February 2014

Keywords:
Network virtualization
Distributed systems
Resource discovery
Resource allocation
Virtual network embedding

ABSTRACT

The European Union Project Networking innovations Over Virtualized Infrastructures (NOVI) set out to design and implement a modular data, control and management plane federation architecture, leading to an integrated experimental prototype mounted on interconnected European Future Internet testbeds. In this paper we present the components of this architecture, responsible for resource discovery and mapping of virtual topologies over a federated multi-domain network virtualization environment. We subsequently introduce a method for the efficient mapping of user requests for virtual networks onto a substrate infrastructure, adopting a semantic-based approach to address the problem. The efficiency of the proposed scheme is evaluated via simulation and critically compared against common non-semantic-based solutions.

1. Introduction

Over the last decade, network virtualization has been introduced as the means to overcome the ossification of the Internet. Large-scale research experimental infrastructures, providing sharing mechanisms, have been widely deployed, enabling Future Internet (FI) experimental validation in realistic testing environments. What is more, network virtualization is considered an inherent component of a polymorphic ecosystem promoting a flexible architectural FI design.

Two initiatives have emerged as main promoters of research on the FI research experimental infrastructures: the Global Environment for Network Innovations (GENI) in the United States and the Future Internet Research and Experimentation (FIRE) in Europe. The EU Project NOVI [1], Networking innovations Over Virtualized Infrastructures, an experimentally driven research project under the FIRE initiative, was set out to create a blueprint of FI federated infrastructures, by designing and prototyping a service portfolio based on combined virtualized facilities from virtualized infrastructures. NOVI architecture has been prototyped as a proof of concept of the proposed design and deployed over Federated E-infrastructure Dedicated to European Researchers Innovating in Computing network Architectures (FEDERICA) [2] and a private deployment of the PlanetLab [3] infrastructures.

In this paper we present the prototyped components of the NOVI architecture, responsible for resource discovery and mapping over a federated multi-domain Network Virtualization

* Corresponding author. Tel.: +30 2107721247.
E-mail addresses: c.pittaras@uva.nl (C. Pittaras), chrisap@noc.ntua.gr (C. Papagianni), arisleiv@netmode.ece.ntua.gr (A. Leivadeas), p.grosso@uva.nl (P. Grosso), vdham@uva.nl (J. van der Ham), papavass@mail.ntua.gr (S. Papavassiliou).

http://dx.doi.org/10.1016/j.future.2014.01.003
0167-739X/

Environment (NVE). On one hand, the problem of mapping virtual networks (VNs) onto a (multi-domain) substrate network is the main resource allocation challenge in network virtualization, commonly referred to as (inter-domain) *Virtual Network Embedding* (VNE) problem or *Virtual Network Mapping* problem. On the other hand resource discovery in the context of an NVE remains still an unexplored research area [4], despite its large impact on the subsequent resource allocation phase. NOVI introduces a distributed, semantic-based resource discovery and VN mapping framework over the NOVI-federated testbeds. The framework is enabled by the NOVI Information Model (NOVI IM) [5] that provides abstractions and semantics of federated virtualized platforms. Semantic technologies facilitate reasoning when selecting resources and services that improves the precision of resource discovery. The semantic-based VNE approach is presented in detail in the study. Due to the small scale of the substrate testbeds used in NOVI, its efficiency is illustrated in a simulation environment that allows for a flexible and comparative performance evaluation.

The rest of the paper is organized as follows. Section 2 summarizes work on resource discovery and allocation in the context of a virtualized environment. Section 3 provides a high level description of the NOVI architecture, identifies the components responsible for distributed resource discovery and mapping and provides a thorough description of the proposed resource discovery and allocation framework. The semantic-based VNE approach is described in Section 4 and is compared against non-semantic aware heuristics, provided by the authors in [6], in Section 5. Finally, Section 6 concludes the paper.

2. Related work

Inter-domain VNE can be basically broken down to the following sub-problems [7]: (i) selecting the appropriate testbed to embed a segment or the entire VN request (VN partitioning) and (ii) solving the resulting distinct VNE problems within each single administrative domain. VN partitioning requires that testbeds advertise their resources to an appropriate resource discovery framework and the VN request graph is split among testbeds (sub) optimally based on advertised costs.

2.1. Resource discovery and VN partitioning

Publications focusing on resource discovery for multi-domain virtualized infrastructures are sparse [4]. A resource discovery method based on the conceptual clustering technique has been proposed [8]. Substrate resources are described in terms of functional and non-functional attributes. Functional attributes define properties of the resources e.g., resource type, Operating System (OS) etc., while non-functional attributes specify criteria and constraints including dynamic features like node and link capacity. Authors in [9] use the same framework as in [8] but propose a hierarchical organization framework of resource discovery data.

In most studies, advertised resource provisioning costs are randomly generated [7,10]. The cost of links spanning multiple domains is assumed to be an order of magnitude higher than the average intra-domain cost. However, dynamic information about resources is not included due to the induced real time monitoring overhead. Belbekkouche et al. [4] pinpoint the necessity to make information on non-functional attributes available at the Service Provider (SP) level. In the context of a networked cloud environment (NCE), authors in [11] defined advertised costs based on resources (VM/network) availability within the networked cloud. The cost of an inter-cloud virtual link is defined as a function of the transit network providers it traverses. Following a similar approach, in the previous work done by the authors [12], advertised costs include resource availability, in terms of scarcity of the resource and the average utilization over a time window.

NOVI follows a semantic-based approach for resource discovery. NOVI's resource discovery framework operates on a distributed set of semantic repositories located on the NOVI service layer of every testbed in the federation [13]. Advertised costs include functional and non-functional attributes, incorporating monitoring mechanisms to acquire aggregated information about each virtualized platform [12].

The problem of VN graph partitioning among several virtualized infrastructures is NP-Hard [7]. VN partitioning has been presented in the literature within the context of an NCE [11,12] or an NVE [7,9]. Within the context of NOVI, the Iterated Local Search meta-heuristic has been adopted and prototyped, due to its intrinsic simplicity and documented efficiency in overcoming time performance/scalability issues [12].

2.2. Virtual network embedding

The problem of assigning virtual nodes and link to a substrate network without violating capacity constraints within a single domain can be reduced to the NP-hard multi-way separator problem [14]. The VN embedding problem is quite challenging, due to finite node and link resource constraints, admission control, and the on-line nature of VN requests. The problem remains computationally intractable even if some conditions are relaxed (i.e., all the requests are known in advance). Several approaches have been followed to deal with the complexity and challenges related to the VNE problem. Extensive literature reviews are provided in the following studies [15,16,4], including comprehensive overviews of the main challenges and diverse aspects of VNE, highlighting existing approaches and emerging requirements.

The NOVI service layer abstracts the physical substrate to a semantic graph. Based on the assumption that experimenters describe the attributes of their resource requirements explicitly, a semantic-based resource mapping approach is followed in NOVI where SPARQL (SPARQL Protocol and RDF Query Language) endpoints enable fast and easy access to substrate resource information.

2.3. Resource discovery and allocation in other projects

The GENI and the FIRE initiatives include a number of projects that face the challenge of resource discovery and allocation. GENI encompasses a number of control frameworks and testbeds e.g., Emulab, ProtoGENI, PlanetLab, Open Resource Control Architecture (ORCA), etc. Slice-based Federation Architecture (SFA) [17] is used as a common API, providing a common method for advertising resources and servicing user requests over multiple administrative domains. The central data structure used by SFA is the Resource Specification (RSpec) [18]. It is used as an interchange format for GENI platforms to advertise substrate resources or describe allocated resources via an appropriate manifest. In addition it enables the user to request resources by providing a complete mapping between requested and substrate resources. However, the ORCA [19] framework is capable of supporting mapping of requested to substrate resources [11]. Specifically topology partitioning is enabled by integrating a minimum k-cut algorithm followed by sub-graph isomorphism. Virtual topology embedding within a single cloud site is facilitated by NEuca [19]. ProtoGENI also defines a slice embedding service without however providing additional details on the service implementation [20]. The slice-based federation architecture is the most common also in the FIRE initiative for the discovery, reservation and provisioning tasks. Within FIRE the use of semantic-based resource descriptions is considered for federated testbeds along with exploiting the expressiveness of semantic technologies for VN embedding [21].