INTERNATIONAL JOURNAL OF CRITICAL INFRASTRUCTURE PROTECTION 6 (2013) 159–178

Available online at www.sciencedirect.com

ScienceDirect

www.elsevier.com/locate/ijcip

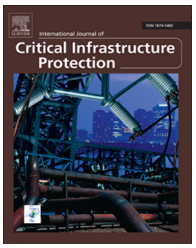

Towards end-to-end network resilience

Panagiotis Vlacheas[a,*], *Vera Stavroulaki*[a], *Panagiotis Demestichas*[a], *Scott Cadzow*[b], *Demosthenes Ikonomou*[c], *Slawomir Gorniak*[c]

[a]*University of Piraeus, 80 Karaoli and Dimitriou Street, 18534 Piraeus, Greece*
[b]*Cadzow Communications Consulting, 10 Yewlands, Sawbridgeworth, Hertfordshire CM21 9NP, United Kingdom*
[c]*European Network and Information Security Agency (ENISA), 1 Vasilissis Sofias, 15124 Marousi, Attica, Greece*

ARTICLE INFO

Article history:
Received 10 November 2011
Received in revised form
16 August 2013
Accepted 16 August 2013
Available online 22 August 2013

Keywords:
Future networks
Network resilience
Threats
Cognitive framework
Ontology
Profiles
Policies

ABSTRACT

Telecommunications networks have evolved towards a unified, service-oriented, operator-governed and autonomic managed infrastructure. Unification ensures interoperability and federation among different domains, technologies, architectures, while allowing the joint consideration of network and service aspects towards a "network as a service" view. Autonomicity reduces operational expenditures and governance guarantees operator control over the entire network. In this new environment, the meaning of network resilience must be revised in an end-to-end manner. This paper focuses on network resilience, identifies the principal network resilience concepts and proposes an ontology, which describes the content and the interactions between the resilience concepts.

1. Introduction

The notion of resilience was introduced in the early decades of the twentieth century in a variety of scientific domains, such as physics, psychology and psychiatry, ecology, business, industrial safety and telecommunications [21]. However, network resilience has recently gained much attention as the property of a network to sustain its normal operations and desired performance when facing a number of predicable or unpredictable situations such as threats and changes.

There have been considerable efforts in the literature (see, e.g., [21]) to distinguish resilience from terms such as fault-tolerance, dependability, robustness, and to determine the boundaries and relationships between resilience and other terms such as stability and diversity. Often, resilience as a more global concept seems to encompass other terms (e.g., dependability) and resilience has even been built into their initial definitions in an incremental approach. After a careful investigation of resilience, it is safe to conclude that resilience evolves in parallel with network development and operating requirements and challenges. For example, fault-tolerance has been known to exhibit some robustness with respect to fault and error handling; and dependability has been used in ubiquitous systems to describe the ability to deliver services that can be trusted in the face of continuous changes [1]. Thus, fault-tolerance and dependability may be considered to be resilience properties.

Trends and challenges regarding future networks impose new requirements when considering resilience. In order to avoid starting from scratch and recognizing the fact that there are already some initiatives that address aspects of future networks, this effort has capitalized on existing

*Corresponding author. Tel.: +30 2104142749.
E-mail address: panvlah@unipi.gr (P. Vlacheas).

1874-5482/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ijcip.2013.08.004

research results. For example, the UniverSelf Project [36] provides a set of top-level requirements and design goals that have to be covered in order to address the general vision and research directions with regard to future networks, service-oriented computing and networking, and the future Internet. These requirements include governance, unification, service orientation, autonomicity, orchestration and coordination and intelligence embodiment. Governance allows a network operator to have control even over an autonomic network by setting goals and requests, enforcing them on the corresponding network elements and receiving notifications when a situation requires prompt operator (re)action. Unification implies interoperability and federation among diverse management systems, which may involve different network segments (radio access, core and service) and may be implemented on different autonomic architectures. Service orientation denotes a joint service and network management in the sense that everything in the network may be considered to be a managed service. Autonomicity allows network entities to be managed and operate with "self" properties such as self-configuration, self-monitoring, self-optimization, self-healing, self-diagnosis, self-protection, self-awareness and self-testing, and, thus, may be considered to be a synonym of self-management. Orchestration and coordination guarantee that simultaneously operating and even conflicting autonomic entities will not cause any instabilities or incompatibilities. Finally, intelligence embodiment represents the progressive introduction of autonomic features in the management chain, and especially in network and service domains, in a distributed manner.

These requirements can guide the redefinition of resilience for future networks. A well-established way to do this is through the design and use of an ontology [37]. An ontology is the term used to refer to the specification of a shared conceptualization within a domain of interest. It involves the definition of domain concepts (e.g., objects, attributes and processes) and their properties and relationships. Therefore, it can be used to model and/or describe the domain, reason about the included entities, and create a unifying framework to solve problems in the domain. Usually, this specification is formal and standardized (explicit ontology), but it may also involve subjective usage (implicit ontology). Ontologies are useful tools for representing knowledge in many scientific domains, such as communications, system and software engineering, enterprise modeling, information architecture, and in general, wherever there is a strong need for a shared vocabulary and interoperability.

Conventional networks comprise several management domains in which resilience is of utmost importance. The management processes and systems in these domains are heterogeneous and typically adhere to specific standards. For the sake of argument, a resilience ontology may be at least applicable to one domain, allowing the definition of the semantic aspects of the domain and its information. In order to address the requirements of future networks, an end-to-end resilience ontology, which integrates information that currently belongs to each domain, is required. Such a "meta-ontology" would allow network managers to elaborate and reason about resilience aspects with an abstract and interoperable view. Moreover, the ability of an ontology to model behavior, in terms of rules and constraints, enables managers to compensate for problematic situations in an autonomic as well as domain-independent manner.

Ontology-based network management [23] has being applied to a variety of network management and security scenarios. A typical example is autonomic management on behalf of an operator of a service deployment constructed on top of a heterogeneous network infrastructure. The proposed end-to-end network resilience ontology will help to hide the heterogeneity of the underlying infrastructure and to formulate an appropriate information flow, while deploying the new services with concrete attributes such as quality of service (QoS) and quality of experience (QoE). Another example is network security and policy management, which are

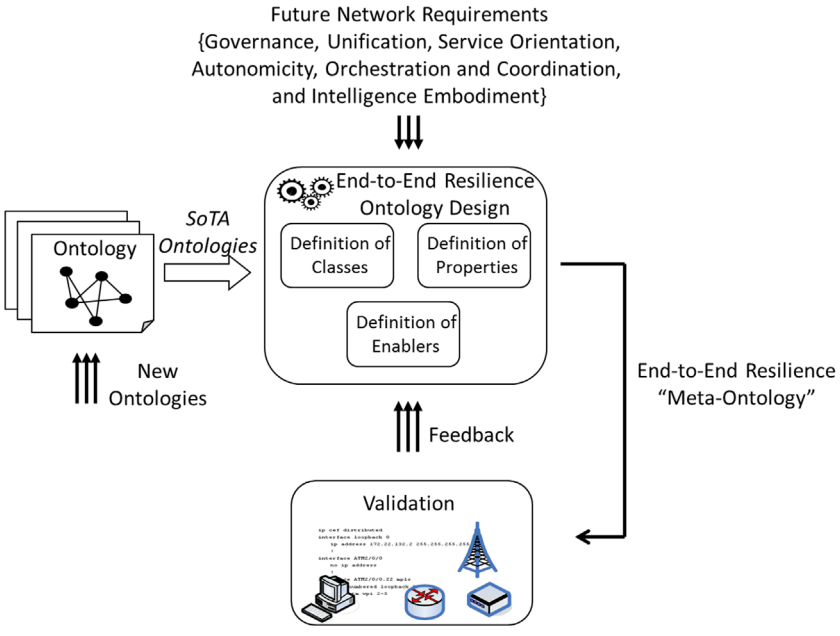

Fig. 1 – Ontology design methodology.