Journal of Network and Computer Applications 36 (2013) 1405–1417

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

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

Automated context dissemination for autonomic collaborative networks through semantic subscription filter generation

Steven Latré [a,*], Jeroen Famaey [a], John Strassner [b], Filip De Turck [a]

[a] Ghent University, Department of Information Technology, IBBT, Belgium
[b] Software R&D Laboratory, Huawei, USA

ARTICLE INFO

Article history:
Received 14 August 2012
Received in revised form
7 December 2012
Accepted 13 January 2013
Available online 5 February 2013

Keywords:
Context dissemination
Autonomic collaborative networks
Future Internet
Semantic approach
Service management
OWL
RDF

ABSTRACT

The current manual management of services and applications in today's telecommunication networks is becoming increasingly complicated. In the Future Internet, management is assumed to be automated by introducing an autonomic layer of distributed management elements. These distributed management elements need to collaborate with each other to ensure end-to-end quality guarantees. In this paper, we focus on the context dissemination between such collaborative management elements. Context dissemination is the exchange of all relevant management data and knowledge between the elements. Collaborating elements typically generate large amounts of context and it is important to filter this continuous stream. We propose a context dissemination approach that automates the context exchange between elements. The approach enables the automated generation of semantic subscription filters. Subscription filters allow an element to define where, how, and when context needs to be requested from other entities. Moreover, the proposed approach allows making the subscription filter generation dependent on the context. We present algorithms that intelligently filter the knowledge that is stored in the ontology. The results show that the generation of subscription filters can be done in the order of tens of milliseconds.

1. Introduction

In recent years, the Internet has evolved from a best-effort packet forwarding network towards a service-oriented delivery framework that supports rich and complex services and applications. To support their management, delivery guarantees must be provided in terms of Quality of Service and Quality of Experience (QoE). Therefore, it is becoming too costly and complex to continue managing these services and applications manually. In the Future Internet, a more automated management approach is required that allows self-governing the network by introducing an intelligent autonomic layer on top of today's network (Jennings et al., 2007). This autonomic layer features a decision making process that supports an automated management of the network's resources. Given the scale and form of the current Internet, it is not possible to maintain a single decision making entity. Instead, an autonomic network management substrate consists of specialized distributed decision making components, called autonomic elements (AEs).

To ensure end-to-end management of the Future Internet, the different AEs need to collaborate with each other. This collaboration is crucial in achieving a well-performing autonomic management framework: it guarantees that AEs do not contradict each other's decisions. One of the primary challenges in designing a collaborative autonomic network is the efficient dissemination of data. In order to collaborate with each other, AEs require management data from other AEs such as (i) monitoring reports falling outside their authority, (ii) decisions taken by other AEs and (iii) knowledge that was inferred by other AEs. All this management data can be seen as context. More specifically, throughout this paper, we call all management data that is relevant to the decision making process of an AE, context. More specifically, we use the following definition of context from DEN-ng (Strassner, 2003): "The Context of an Entity is a collection of measured and inferred knowledge that describe the state and environment in which an Entity exists or has existed". In particular, our definition emphasizes two types of knowledge: facts (that can be measured) and inferred data, which results from machine learning and reasoning processes applied to past and current context. It also includes context history, so that current decisions based on context may benefit from past decisions, as well as observations of how the environment has changed.

Many autonomic architectures use the publish-subscribe paradigm to communicate with other AEs and enable the collaboration between them. The publish-subscribe paradigm allows consumers of context (i.e., AEs) to express their interest in context by subscribing to that context at the producer (i.e., another AE) (Eugster et al., 2003). This subscription is done through filters,

* Corresponding author. Tel.: +32 93314988.
E-mail address: steven.latre@intec.ugent.be (S. Latré).

1084-8045/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.jnca.2013.01.011

which define what type of context producers need to send to the consumers. Traditionally, these subscription filters are statically defined by the consumer interested in the context. In this paper, we focus on the automated generation of subscription filters between collaborating AEs in the autonomic networking paradigm. We focus specifically on the context dissemination process inside an administrative domain and declare the context dissemination between administrative domains as out of scope. In this context, an administrative domain is a subnetwork that is managed by a single network operator. Inside an administrative domain, the main challenge of context dissemination is its scalability. Large amounts of context are generated by different AEs and the type of context that needs to exchanged can fluctuate rapidly. Instead, between administrative domains, other challenges are notable such as the negotiation of the context dissemination: this is part of future work. We propose a context dissemination process that allows coping with the fluctuation of context requirements by generating the subscription filters automatically. The process takes into account the requirements and goals of a specific AE in terms of context, as well as changes in the state of the environment.

The remainder of this paper is structured as follows: an illustrative use case of context dissemination for multimedia delivery in an access network is described in Section 2. Section 3 discusses similar work in the field of context modeling in network management. The context dissemination process is described in full in Sections 4 and 5, which focus on the design of the semantic model and the algorithmic contributions, respectively. Finally, Section 6 presents detailed evaluation results of the performance of the proposed approach.

2. Use case: autonomic multimedia service management in access networks

To provide a better understanding of the context dissemination process, we describe it for an illustrative use case. This use case is centered around the autonomic management of multimedia services in access networks. We assume that a service provider is offering a number of multimedia services to the end users. These multimedia services consist of the typical video-based services (e.g., digital TV broadcasting, Video on Demand) that are part of today's triple play offers in IPTV environments. As illustrated in Fig. 1, the content originates at the video head-end in the access network. The service provider can use its own access network, consisting of caches, access routers and access nodes to stream the videos to its customers, residing in the home network. The streaming is done to set-top boxes (STBs), which are devices located in the home network that are capable of playing the video on a television screen. Despite being located in the home network, the service manager is typically the owner of a STB and therefore has access to it for management purposes.

To manage the multimedia services, a set of collaborating AEs are logically deployed on top of the access network. Each AE governs a specific type of device. The AEs collaborate with each other by (i) exchanging context with each other and (ii) requesting management actions from one AE to another, which can then propagate this request to the device(s) it manages. For example, every STB has its own AE that can monitor the delivered quality of the multimedia streams and offers an interface to remote AEs to make configuration changes (e.g., a request to send more context or a request to alter a configuration parameter at the video client).

In this use case, we focus on the context dissemination process of the video head-end AE. We assume that the video head-end AE has a number of management tasks available to manage the multimedia services. For example, it can reduce the streamed video quality if congestion occurs in the network. Alternatively, it can protect the video against packet loss due to lossy links by adding additional redundancy. In order to determine the optimal configuration of these management tasks, it requires knowledge about the status of the managed network. However, at the same time, it is useless to continuously request all this knowledge if the current management configuration is streaming the video at a

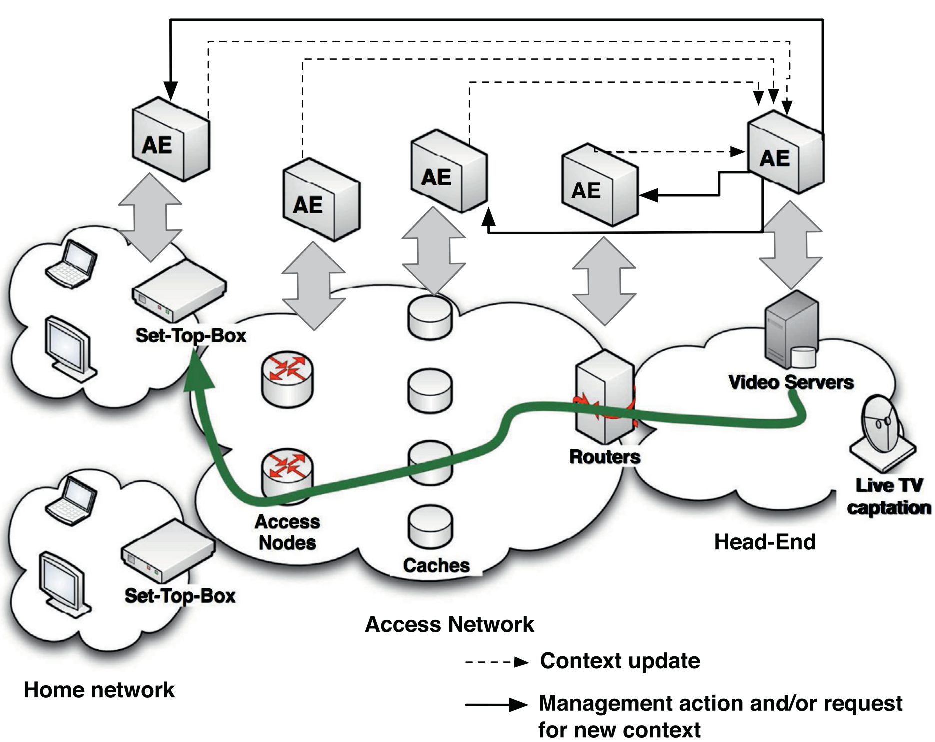

Fig. 1. Overview of the illustrative use case, representing a multimedia delivery scenario in an access network. Different autonomic elements (AEs) exchange context with each other to maintain the highest possible video quality.