Future Generation Computer Systems 56 (2016) 77–94

Contents lists available at ScienceDirect

Future Generation Computer Systems

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

A trust-aware, self-organizing system for large-scale federations of utility computing infrastructures

F. Messina [a], G. Pappalardo [a], D. Rosaci [b,*], C. Santoro [a], G.M.L. Sarné [c]

[a] *Department of Computer Science and Mathematics, University of Catania, Viale A.Doria 6, 95126 Catania CT, Italy*
[b] *Department DIIES, University "Mediterranea" of Reggio Calabria, Loc. Feo di Vito, 89122 Reggio Cal. RC, Italy*
[c] *Department DICEAM, University "Mediterranea" of Reggio Calabria, Loc. Feo di Vito, 89122 Reggio Cal. RC, Italy*

HIGHLIGHTS

- We discuss a trust-based, decentralized approach for Utility Computing Federations.
- Our approach is suitable for federations in which partners are in competition.
- We designed an overlay network for fully decentralized, efficient finding approach.
- We also show that the trust model allows to find the most promising interlocutor.

ARTICLE INFO

Article history:
Received 8 January 2015
Received in revised form
23 May 2015
Accepted 25 July 2015
Available online 7 September 2015

Keywords:
Distributed systems
Application and expert systems
Information systems application

ABSTRACT

On-demand computing environments, like Cloud/Grid systems, consist of nodes that individually manage local resources intended to be served to clients. When a client needs some resources, it has the problem of finding the most suitable nodes capable of providing them. In addition, a provider node too may be in need to efficiently locate resources for itself, given the emerging, highly *competitive*, context of large-scale federations. Indeed, a node competes, with the other federated ones, to obtain the assignment of available tasks. To this end, it may decide to publish a set of resources/ services wider than the one it has currently available. Should such a node be assigned a job for which its actual resources are insufficient, it could end up requiring the collaboration of other nodes.

Hence the crucial problem, for nodes and clients alike, is to determine the most promising collaborators. For this purpose, in the competitive and demanding scenarios considered, we advocate taking into account the *trustworthiness* of nodes in declaring their capabilities, i.e., to help it making an effective selection of possible collaborators, each node should be provided with a trust model for accurately evaluating the trustworthiness of its interlocutors.

In this paper, a trust-based approach for large-scale federations Utility Computing infrastructures is proposed. The proposed model is designed to allow any node to find the most suitable collaborators in an efficient way, avoiding exploration of the whole node space. A fully decentralized approach is employed, which allows nodes of a federation to be organized in an overlay network on the basis of suitable criteria. This enables any customer or provider in need of collaborators to determine a suitable set of candidate nodes within which to search.

1. Introduction

Since the development of computational Grids [1], started in the middle of 1990*s*, virtualization technology [2,3] and commodity hardware led to the development of the Cloud computational paradigm [4,5], which is a practical example of a new distributed computing paradigm, i.e. *On-demand computing* [5]. The next step is represented by the construction of federated computing infrastructures, i.e. Grid Federations [6] and federated Cloud Computing infrastructures [7–10]. For instance, we can cite the Reservoir model developed by IBM [11], in order to "*facilitate an open, service-based online economy in which resources and services are transparently provisioned and managed across clouds on an on-demand basis*

* Corresponding author.
E-mail addresses: messina@dmi.unict.it (F. Messina), pappalardo@dmi.unict.it (G. Pappalardo), Rosaci@unirc.it (D. Rosaci), santoro@dmi.unict.it (C. Santoro), sarne@unirc.it (G.M.L. Sarné).

http://dx.doi.org/10.1016/j.future.2015.07.013
0167-739X/

at competitive costs with high-quality service", bringing a vision of "*implementing an architecture that would enable providers of cloud infrastructure to dynamically partner with each other [...] while fully preserving their individual autonomy in making technological and business management decisions*".

The increasing interest for large-scale federations of utility computing infrastructures is motivated by the opportunity for providers to mutually "rent" virtualized services and, in turn, allocate unused resources in order to increase profits [12]. That is one of the main reasons to have interest in federating resources: establishing a *collaborative* context at large-scale will help to achieve the above objectives. For instance, providers of computing infrastructures (IaaS, PaaS, STaaS) are able to provide services like *auto-scaling, advanced monitoring, automation in service deployment* to providers which make their business by offering high level e-services (e.g. SaaS, DaaS) [13]. In such a collaborative scenario, *fulfilling computational requirements* of user requests *as efficiently as possible* becomes crucial: this is a typical problem in a large-scale distributed system, in which resources are shared by several different organizations [14]. For instance, a recent study [15] deals with the problem of balancing profit, price and QoS in a typical federated environment, by proposing policies to help making decisions aimed at outsourcing services, contributing to the federation or even terminating spot VMs to reject less profitable requests.

Another important issue to address is that service providers – while having interest in joining a federated network that generally results more attractive for users than a single provider – are in *competition* with each other, in order to offer a number of *interchangeable* on-demand services at any level. The increasing trend of offering interchangeable services by federated providers leads to the interesting recent development [16] of Cloud Service Composition, a popular choice affording clients the opportunity of choosing among several alternatives. Yet, in a highly competitive environment, this opportunity carries a risk, in that nodes *can declare to provide a better QoS than currently possible*, as providers have a "natural" tendency to gamble with their own *reputation* [12]. Since Cloud customers tend to prefer paying services in the face of signed SLAs (Service Level Agreements) [17,18], when negotiated QoS is not fulfilled, both the node provider and the user are penalized. While the former will generally pay a refund [17], the activities of the latter (the customer) will be affected by a measured level of QoS which is not compliant with the SLA, which may lead to a loss of revenue [13] uncompensated by the refund. In order to avoid penalties and losing reputation with clients, a node, unable to fulfill an assigned client request, may seek collaboration from other nodes at a price, effectively outsourcing, in turn, some of the services it should provide. However, in so doing, it may incur the hazard of not finding collaborators that are able to provide a sufficient level of reliability themselves. In this scenario, an accurate selection of interlocutors, aimed at *finding the most promising collaborators*, can make a great difference. As we discuss below, this is a major issue we deal with in this work, i.e. to provide a decentralized support to assess the trustworthiness of Cloud providers in a federated context.

Given these premises, we propose and discuss a decentralized approach, called *Hypertrust*, for federated, large scale utility computing infrastructures, aimed at providing a twofold support to customers and providers. First of all, Hypertrust deals with trustworthiness in a competitive, large-scale distributed system. The basic idea is to *design a trust system* suitable to support the nodes in choosing their interlocutors, *limiting the search space* to only those nodes declaring to have the suitable resources for the collaboration. The second feature is motivated by the fact that, in spite of existing trust models like RRAF [19,20] were introduced to combine several different concerns related to trustworthiness of computational nodes into a unique synthetic trust measure, they demonstrate good performance when applied to a small/medium-size system and, when the system includes a very large number of nodes, as in the case of large-scale federated computing systems, this selection task becomes infeasible.

The last target, i.e. *fulfilling computational requirements of user requests as efficiently as possible* is pursued through a *decentralized* technique, intended to organize computing nodes as peers of an overlay network whose "small world-like" characteristics make the resource finding process both effective and efficient. Indeed, the overlay construction algorithm is guided by *resource status similarity*, so that peers featuring a similar amount of resource availability tend to be interconnected (in terms of overlay links), thus forming clusters which, in turn, are connected together though sparse "long" links.[1] Not surprisingly, resource finding by navigating such an overlay network, turns out to be rather efficient. This is how we address the first issue of (*efficient resource finding*).

Some special nodes are also introduced, playing the role of *Task Allocators* (*TA*): they are mediators, provided in principle by each organization, which collect users' feedbacks and, based on these, select a suitable node for a service. As TAs collect feedbacks about service requests for which they have been contacted, it makes sense for users to generally rely on them, thus exploiting previous user experiences. In contrast, a computational node, in need to select backup collaborators to provide a service, will resort to our decentralized resource finding algorithm to determine a set of candidate nodes, and then exploit the trust model, to choose a collaborator.

We simulated some experimental scenarios in order to highlight the advantages introduced by our proposed approach. Results prove that, in a competitive scenario, recourse to the Hypertrust overlay network, enhanced with the trust model for interlocutor selection, affords a significant advantage with respect to selection based on resource declaration alone. Moreover, as discussed in Section 5.6, not only do TA components bridge the gap between the two models performing the selection needed by the Hypertrust selection system: they also bring a further contribution, in terms of QoS, to the final user. We have also performed an experimental analysis of our decentralized resource finding algorithm, employed to seek collaborators with suitable characteristics. Aptly, it has shown quite a high efficiency, despite the lack of any centralized directory or index.

The remaining of the paper is organized as follows. Section 2 introduces the competitive, large-scale distributed scenario we deal with. Section 3 describes the adopted trust model and brings out a brief discussion about existing models. In Section 4, the Hypertrust overlay construction and resource finding algorithm are discussed, with particular emphasis on Cloud Federations. In Section 5 we provide an experimental evaluation of the proposed approach. Related work is presented in Section 6. Finally, Section 7 reports our conclusions and discusses possible developments of our ongoing research.

2. The competitive hypertrust scenario

An increasing attention is emerging about the dynamics of resource pricing in multi-cloud/grid contexts where, similarly to other competitive utility markets, each provider node is driven by profits, while their clients seek lower costs and/or higher quality-of-service.

Several works deal with this topic, but some of them completely ignore the competitive nature of distributed computing

[1] Such an organization resembles the classical model of small-world networks [21].