Computers and Electrical Engineering 000 (2017) 1–14

Contents lists available at ScienceDirect

Computers and Electrical Engineering

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

Performance analysis of multi-institutional data sharing in the Clouds4Coordination system☆

Ioan Petri[a], Omer F. Rana[a,*], Tom Beach[b], Yacine Rezgui[b]

[a] School of Computer Science & Informatics, Cardiff University, Wales, United Kingdom
[b] School of Engineering, Cardiff University, United Kingdom

ARTICLE INFO

Article history:
Received 22 February 2016
Revised 11 February 2017
Accepted 13 February 2017
Available online xxx

Keywords:
Cloud computing
Cloud federation
Performance analysis
Distributed coordination
CometCloud
Architecture
Engineering and construction (AEC) projects

ABSTRACT

Cloud computing is used extensively in Architecture/ Engineering/ Construction projects for storing data and running simulations on building models (e.g. energy efficiency/environmental impact). With the emergence of *multi-Clouds* it has become possible to link such systems and create a distributed cloud environment. A multi-Cloud environment enables each organisation involved in a collaborative project to maintain its own computational infrastructure/ system (with the associated data), and not have to migrate to a single cloud environment. Such infrastructure becomes efficacious when multiple individuals and organisations work collaboratively, enabling each individual/ organisation to select a computational infrastructure that most closely matches its requirements. We describe the "Clouds-for-Coordination" system, and provide a use case to demonstrate how such a system can be used in practice. A performance analysis is carried out to demonstrate how effective such a multi-Cloud system can be, reporting "aggregated-time-to-complete" metric over a number of different scenarios.

1. Introduction

Cloud computing enables applications in the Architecture/ Engineering/ Construction (AEC) to dynamically scale-up (increasing volume of data/computation) & scale-out (increasing diversity of computational infrastructures involved). This becomes more relevant when projects are being undertaken by a consortia of companies, who work collaboratively for the duration of the project. Such projects are complex and the consortia members provide a range of skills to the project from its inception to completion. During this process, various data artifacts are also generated that need to be stored and shared between project members (generally using access control strategies – which limit what can be accessed at a particular stage of the AEC project lifecycle). The planning, implementation and running of these AEC industry projects requires the formation of secure Virtual Enterprises (VEs) to enable collaboration between its members by sharing project information and resources. An important feature of the consortia is that they are dynamic in nature and are formed for the lifetime of the project. Members can participate in several consortia at the same time and can join or leave a consortium as the project evolves.

☆ Reviews processed and approved for publication by Editor-in-Chief.

* Corresponding author.
E-mail addresses: petrii@cardiff.ac.uk (I. Petri), ranaof@cardiff.ac.uk (O.F. Rana), beachth@cardiff.ac.uk (T. Beach), rezguiy@Cardiff.ac.uk (Y. Rezgui).

http://dx.doi.org/10.1016/j.compeleceng.2017.02.015

Cloud computing offers an important computing infrastructure to facilitate the establishment and coordination of such VEs. As well as remote access, Cloud computing also provides enhanced security, including single sign-on capability, security between consortia members, simple setting up of networks to support VEs, distribution of computationally intensive jobs across multiple distributed processors (based on shared information about available resources). Each organisation involved in a VE may have access to its own Cloud computing system (privately managed internally within the organisation, or acquired through a public provider such as Amazon.com or Microsoft (via their Azure platform)). It is unlikely that all members of a consortium will share the same platform. Integrating capability across multiple platforms is therefore an essential requirement for such VEs to function in an efficient and reliable manner. The alternative would be for all members of the consortia to migrate to the same platform. Unless there is industry-wide (or consortia-wide) agreement on what this platform should be, such a migration process can be costly and error prone, and often the outcome of the process for a specific project can be unclear.

Various efforts have been proposed to implement such multi-Clouds, ranging from research efforts focused on Cloud interoperability e.g. the Open Cloud Computing Interface (OCCI) efforts at the Open Grid Forum [1]. OCCI provides an API and a set of protocols to enable management capability to be carried out across multiple Cloud providers. A variety of implementations are currently available, in systems such as OpenStack and OpenNebula (two open source Cloud platforms). An alternative approach to interoperability is through the development of specialist gateway nodes which enable mapping between different Cloud systems and the implementation of specialist gateways to connect different Cloud systems, the development of a Cloud Operating System (CloudOS) to connect distributed Clouds to the use of specialist in-network capability to process data in network elements between different end points (GENICloud [2]). Similarly, on-line sites such as CloudHarmony [3] report over 100+ Cloud providers that offer capability ranging from storage and computation to complete application containers that can be acquired at a price, primarily using service-based access models. As the multi-Cloud market and associated number of Cloud providers who could offer services in such a market increase in number, there is often a need to understand which Cloud providers are likely to be of most benefit in the context of a given application requirement. Matchmaking becomes an important capability in such a marketplace – enabling application users to map their requirements to infrastructure capability that may be hosted across a number of different types of Cloud systems.

This paper has two major contributions: (i) we present the implementation and use of a distributed Clouds4Coordination system used to coordinate large construction projects based on requirements of the AEC sector. We emphasise the need to aggregate capability across multiple systems, rather than require all project members to migrate to a single system, and (ii) we demonstrate the capability of the system to provide the required quality of service and functionality in use by running performance analysis and measurement. Our approach involves the implementation of a logical “shared” space that is physically distributed across multiple sites involved in a cloud federation. Such a shared coordination space enables various project members to interact with each other during the stages of a project. We compare our approach to general cloud federation efforts, specifically adapted for the needs of the AEC industry in Section 2. In Section 3 we present the CometCloud system and how this system has been used to create the federated cloud framework, followed by a description of the “Cloud4Coordination” (C4C) system and the associated Application Programming Interface (API) that makes use of CometCloud in Sections 4 and 5. In Section 6 we evaluate the C4C system in terms of its performance with a scenario, and provide overall conclusions in Section 7.

2. Related work

Through the federation of cloud systems it has become possible to connect local infrastructure providers to a common framework where participants can exchange data and collaborate. The mechanisms used to support cloud federation can bring substantial benefits for service providers by offering facilities for accessing global services instead of increasing costs associated with building new infrastructure (which may not be fully utilized and may only be needed to support peaks in workload over short time frames). More importantly, organisations with spare capacity in the data centre are now provided with a simple way to monetize that capacity by submitting it to the marketplace for other providers to buy, creating an additional source of revenue. Even if computational infrastructure was made available, it may not be possible to host services or data due to issues associated with licensing and intellectual property. Federation in cloud systems has led to a real *democratisation* of cloud markets – enabling businesses to make use of a variety of different cloud providers in different geographic areas. A federated cloud also enables users to host applications with their cloud provider of choice – thereby making local decisions about pricing, software libraries/ systems and deployment environments, while still being able to connect to other computational resources.

In a federation context there are several parameters that need to be considered in order to determine the type of interactions possible between sites. When two or more sites come together, it is important to identify not only the incoming workload of each site but also the cost of outsourcing to resources managed externally, the revenue obtained from outsourcing tasks or the cost of maintaining a reasonable level of utilisation. Identifying a set of such parameters is a challenging task due to the variability in the parameters of a federated environment (such as number of resources allocated to local vs. remote jobs, how many jobs to outsource to another site, the time interval over which access to remote jobs should be allowed, etc) and the fluctuation of resource demand. Depending on the value of such parameters, a site manager must decide whether to outsource resources, compute tasks locally or reject remote task requests altogether [6].

Please cite this article as: I. Petri et al., Performance analysis of multi-institutional data sharing in the Clouds4Coordination system, Computers and Electrical Engineering (2017), http://dx.doi.org/10.1016/j.compeleceng.2017.02.015