

Multi-agent System for Controlling a Cloud Computing Environment

Fernando de la Prieta¹, María Navarro¹, Jose A. García¹,
Roberto González², and Sara Rodríguez¹

¹ University of Salamanca, Computer Science and Automatic Control Department
Plaza de la Merced s/n, 37007, Salamanca, Spain

{fer,mar90ali94,rgonzalezramos,srg,jalberto}@usal.es

² Centro de Innovación Tecnológica CENIT, Salamanca, Spain
jgarccor@insags.com

Abstract. Nowadays, a number of computing paradigms have been proposed, of which the latest one is known as Cloud computing. Cloud computing is revolutionizing the services provided through the Internet, and is continually adapting itself in order to maintain the quality of its services. In this paper it proposes a cloud platform for storing information and files by following the cloud paradigm. Moreover, a cloud-based application has been developed to validate the services provided by the platform.

Keywords: Cloud Computing, multi-agent system, agent-based cloud computing, cloud storage, utility computing.

1 Introduction

The term “Cloud Computing” defined the infrastructure as a “Cloud” from which businesses and users are able to access applications from anywhere in the world on demand. Thus, the computing world is rapidly transforming towards developing software for millions to consume as a service, rather than to run on their individual computers. As a result, the number of both closed and open source platforms has been rapidly increasing [2]. Although at first glance this may appear to be simply a technological paradigm, reality shows that the rapid progression of Cloud Computing is primarily motivated by economic interests that surround its purely computational or technological characteristics [1]. Since user requirements for cloud services are varied, service providers have to ensure that they can be flexible in their service delivery while keeping the users isolated from the underlying infrastructure.

Nowadays, the latest paradigm to emerge is that of Cloud computing which promises reliable services delivered through next-generation data centers that are built on virtualized compute and storage technologies. Cloud computing platforms has properties of clusters or grids environments, with its own special attributes and capabilities such strong support for virtualization, dynamically composable services with Web Service interfaces, value added services by building on Cloud compute, application

services and storage. On this last point, it is important to note that information storage is not performed in the same way today as it was in the past. During the incipient stages of computer sciences, information was stored and accessed locally in computers. The storage process was performed in different ways: in data files, or through the use of database management systems that simplified the storage, retrieval and organization of information, and were able to create a relationship among the data. Subsequently, data began to be stored remotely, requiring applications to access the data in order to distribute system functions; database system managers facilitated this task since they could access data remotely through a computer network. Nevertheless, this method had some drawbacks, notably that the users had to be aware of where the data were stored, and how they were organized. Consequently, there arose a need to create systems to facilitate information access and management without knowing the place or manner in which the information was stored, in order to best integrate information provided by different systems.

This paper presents a Cloud architecture developed in the +Cloud system [5] to manage information. +Cloud is a Cloud platform that makes it possible to easily develop applications in a cloud. Information access is achieved through the use of REST services, which is completely transparent for the installed infrastructure applications that support the data storage. In order to describe the stored information and facilitate searches, APIs are used to describe information, making it possible to search and interact with different sources of information very simply without knowing the relational database structure and without losing the functionality that they provide. Finally, in order to validate the functionality of the services proposed by the +Cloud platform, the paper presents Warehouse 3.0, a cloud storage application.

This paper is structured as follows: the next section provides an overview of the +Cloud platform; Section 3 presents the cloud-based application developed to validate the services provided by the platform: Warehouse 3.0; and finally some conclusions are shown.

2 +Cloud Platform

The platform is composed by a layered structure that coincides with the widely accepted layered view of cloud-computing [3]. This platform allows services to be offered at the PaaS (Platform as a Service) and SaaS (Software as a Service) levels.

The SaaS (Software as a Service) layer is composed of the management applications for the environment (control of users, installed applications, etc.), and other more general third party applications that use the services from the PaaS (Platform as a Service) layer. At this level, each user has a personalized virtual desktop from which they have access to their applications in the Cloud environment and to a personally configured area as well. The next section presents the characteristics and modules of PaaS Layer in +Cloud and +Cloud in greater detail. Both the PaaS and SaaS layers are deployed using the internal layer of the platform, which provides a virtual hosting service with automatic scaling and functions for balancing workload. Therefore, this platform does not offer an IaaS (Infrastructure as a Service) layer.

The virtual and physical resources are managed dynamically. To this end, a virtual organisation of intelligent agents, that is off the topic covered in this paper, and that monitor and manage the platform resources is used [4][5].

The *PaaS Layer* provides its services as APIs, offered in the form of REST web services. The most notable services among the APIs are the identification of users and applications, a simple non-relational database, and a file storage service that provides version control capabilities and emulates a folder-based structure.

The services of the Platform layer are presented in the form of stateless web services. The data format used for communication is JSON, which is more easily readable than XML and includes enough expression capability for the present case. JSON is a widely accepted format, and a number of parsing libraries are available for different programming languages. These libraries make it possible to serialize and de-serialize objects to and from JSON, thus facilitating/simplifying the usage of the JSON-based APIs.

The *FSS (File Storage Service)* provides an interface to a file container by emulating a directory-based structure in which the files are stored with a set of metadata, thus facilitating retrieval, indexing, searching, etc. The simulation of a directory structure allows application developers to interact with the service as they would with a physical file system. A simple mechanism for file versioning is provided. If version control is enabled and an existing file path is overwritten with another file, the first file is not erased but a new version is generated. Similarly, an erased file can be retrieved using the “restore” function of the API. In addition to being organized hierarchically, files can be organized with taxonomies using text tags, which facilitates the semantic search for information and makes the service more efficient. The following information is stored for each file present in the system: (i) Its virtual path as a complete name and a reference to the parent directory. (ii) Its length or size in bytes. (iii) An array of tags to organize the information semantically. (iv) A set of metadata. (v) Its md5 sum to confirm correct transfers and detect equality between versions. (vi) Its previous versions.

Web services are implemented using the web application framework Tornado¹ for Python. While Python provides excellent maintenance and fast-development capabilities, it falls short for intensive I/O operations. In order to keep file uploads and downloads optimized, the APIs rely on the usage of the Nginx² reverse proxy for the actual reads and writes to disk. The actual file content is saved in a distributed file system so that the service can scaled, and the workload is distributed among the front-end servers by a load balancer. The structure of the service allows migrating from one distributed file system to another without affecting the client applications.

File metadata and folder structure are both stored in a MongoDB³ database cluster, which provides adequate scalability and speed capabilities for this application. Web service nodes deploy Tornado and Nginx as well as the distributed file system clients

¹ <http://www.tornadoweb.org/>

² <http://nginx.org/>

³ <http://www.mongodb.org/>

Table 1. Restfull web services exposed by FSS

REST Web Call	Description
PutFile	creates a new file (or a new version of an existing file) in response to a request containing the file and basic metadata (path, name and tags) in JSON, structured in a standard multipart request.
Move	changes the path of a file or a folder
Delete	deletes a file. Can include an option to avoid the future recovery of the file, erasing it permanently
GetFolderContents	returns a JSON array with a list of the immediate children nodes of a specific directory.
GetMetadata	returns the metadata set of a file or directory providing its identifier or full path.
GetVersions	returns the list of all the recoverable versions of a file.
DownloadFile	returns the content of a file (a specific older version can be specified).
Copy	creates a copy of a file or a recursive copy of a folder.
CreateFolder	creates a new folder given its path.
DeleteVersion	permanently deletes a specific version of a file.
Find	returns a list of the children nodes of a folder (recursively).
GetConfiguration	retrieves the value of a configuration parameter for the application.
SetConfiguration	sets the value of a configuration parameter (e.g. enabling or disabling version control)
GetSize	retrieves the size of a file. If a folder path is passed, then the total size of the folder is returned.
RestoreVersion	sets an older version of a file as the newest.
Undelete	restores a file.

(GlusterFS⁴/NFS), and the access to the MongoDB cluster that can be located either within or exterior to the nodes.

The *OSS (Object Storage Service)* is a document-oriented and schemaless database service, which provides both ease of use and flexibility. In this context, a document is a set of keyword-value pairs where the values can also be documents (this is a nested model), or references to other documents (with very weak integrity enforcement). These documents are grouped by collections, in a manner similar to how tuples are grouped by tables in a relational database. Nevertheless, documents are not forced to share the same structure. A common usage pattern is to share a subset of attributes among the collection, as they represent entities of an application model. By not needing to define the set of attributes for the object in each collection, the migration between different versions of the same application and the definition of the relationships among the data become much easier. Adding an extra field to a collection is as easy as sending a document with an extra key. A search on that key would only retrieve objects that contain it. The allowed types of data are limited to the basic types present

⁴ <http://www.gluster.org/>

in JSON documents: strings, numbers, other documents and arrays of any of the previous types.

As with the FSS, the web service is implemented using Python and the Tornado framework. By not managing file downloads or uploads, there is no need to use the reverse proxy that manages them in every node; therefore Nginx is used only to balance the workload at the entry point for the service.

Table 2. Restfull web services exposed by OSS

REST Web Call	Description
Create	creates a new object inside a collection according to the data provided. It returns the created object, adding the newly generated identifier. If the collection does not exist, it is created instantly.
Retrieve	retrieves all objects that match the given query.
Update	updates an object according to the data provided (the alphanumeric identifier of the object must be provided).
Delete	deletes all objects that match the given query.

The *Identity Manager* is in charge of offering authentication services to both customers and applications. Among the functionalities that it includes are access control to the data stored in the Cloud through user and application authentication and validation. Its main features are:

- Single sign-on web authentication mechanism for users. This service allows the applications to check the identity of the users without implementing the authentication themselves.
- REST calls to authenticate application/users and assign/obtain their roles in the applications within the Cloud.

3 Warehouse

Every user has a root folder that contains all their data. The information stored by a user can be shared with other users through invitations for specific folders or using the user's "Public" folder. It is also possible to create groups of users, which work in a similar way to e-mail groups, in order to allow massive invitations. The user interface is updated asynchronously by using WebSockets. The changes made by a user over a shared resource are automatically displayed in the browsers of the other users. The application has syntactic and semantic search capabilities that are applied to different types of files (text, images or multimedia) due to the extraction and indexing of both the textual content and the metadata present in those files. Furthermore, the search results are presented next to a tag cloud that can be used to refine the searches even more. Finally, the application allows users to retrieve and manipulate different versions of their files. This function is powered by the mechanisms present in the underlying file storage API that has been previously described.




Fig. 1. Snapshot of the user interface

The contents of the files and the file system structure are stored in the FSS. Additional information is necessary to establish relationships between the data and to maintain the folder-sharing logic. This extra information is stored in the OSS. Due to the scalability and high-performance of the APIs, the application can execute tasks that will maintain the referential integrity of its model and the high number of recursive operations that are necessary to move and copy folders. When a user creates a folder, three properties are assigned to it automatically: (i) Host user: keeps permissions over the folder. A number of operations are reserved for this user: move, delete, rename and cancel sharing; (ii) A list of invited users, initially empty; and finally, (iii) A list of users with access privileges to the file, initially containing only the host user.

This tool makes intensive usage of both the file and object storage services and it serves the purpose of being the first real-application test for the developed APIs. Warehouse is the first non-native application that has been developed for the +Cloud platform. Using last-generation standards such as HTML5 and WebSockets⁵, the tool allows storing and sharing information using the cloud environment. The user interface is shown at Fig. 1. The available mechanisms for file uploading include HTML5's drag&drop technique. The sharing algorithms are capable of multi-level folder sharing: a children folder of one shared folder can be shared with another list of users. This second group of users will only be allowed to navigate the most-nested folder. The actions related to folder sharing include:

- Invite: adds a user to the list of invited users.
- Accept or decline invitation: If the invitation is accepted, the user is added to the list of access-allowed users. Otherwise, the target user is removed from the list of invited users.
- Leave folder: the user that leaves the folder is removed from the list of access-allowed users. If the host user leaves the folder, the folder will be moved to

⁵ <http://www.websocket.org/>

another user's space and that user will be the new host. If there is more than one user remaining, the current host must choose which user will be the new host.

- Turn private: this operation can only be executed by the host user, and deleting all invitations and resetting the access list.
- Move: if the host moves the file, the other users will see a change in the reference to the shared folder. If the operation is done by another user, then only the reference of that user is modified (no move operation is performed).
- Delete: only the host can execute this operation. The shared folder can be moved to the space of another user, or be completely removed.

The next figure depicts the layered architecture of the application Warehouse 3.0. There are three layers, interface, control and model: (i) the interface layer is developed using HTML5 and jQuery, (ii) the control layer is developed using WebSocket for automatically updating runtime information among all users Warehouse using the system, and finally, (iii) the persistence layer that implemented a DAO pattern specially developed to manage the persistence of information OSS (Object Storage System) and files FSS (File Storage System) in the Cloud Computing environment.

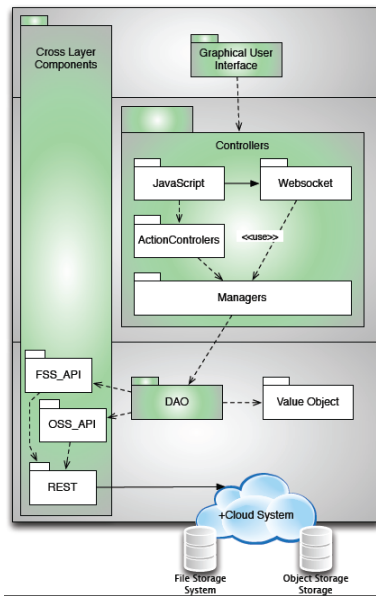


Fig. 2. Layered Architecture Warehouse 3.0

4 Conclusions

Cloud computing is a new and promising paradigm delivering computing services. As Clouds are designed to provide services to external users, providers need to be compensated for sharing their resources and capabilities. In this paper, we have proposed

architecture, +Cloud, that has made it possible to store information in applications without having previously established a data model. The storage and retrieval of information is done transparently for the applications, and the location of the data and the storage methods are completely transparent to the user. JSON can define information that is stored in the architecture, making it possible to perform queries that are more complete than those allowed by other cloud systems. This characteristic makes it possible to change the infrastructure layer of the cloud system, facilitating the scalability and inclusion of new storage systems without affecting the applications. In conclusion, we have presented various Cloud efforts in practice from the service-oriented perspective to reveal its emerging potential for the creation of third-party services to enable the successful adoption of Cloud computing, such as object and file storage infrastructure.

Acknowledgements. This research has been supported by the project *SOCIEDADES HUMANO-AGENTE: INMERSION, ADAPTACION Y SIMULACION*. TIN2012-36586-C03-03.(Ministerio de Ciencia e Innovación. Proyectos de Investigación Fundamental No Orientada). Spain.

References

1. Buyya, R., Yeo, C.S., Venugopal, S.: Market-oriented cloud computing: Vision, hype, and reality for delivering it services as computing utilities. In: 10th IEEE International Conference on High Performance Computing and Communications, HPCC 2008, pp. 5–13. IEEE (September 2008)
2. Peng, J., Zhang, X., Lei, Z., Zhang, B., Zhang, W., Li, Q.: Comparison of several cloud computing platforms. In: 2nd International Symposium on Information Science and Engineering, ISISE 2009, pp. 23–27. IEEE Computer Society (2009)
3. Mell, P., Grance, T.: The Nist Definition of Cloud Computing. In: NIST Special Publication 00-145, pp. 1–3. NIST (2011)
4. Heras, S., De la Prieta, F., Julian, V., Rodríguez, S., Botti, V., Bajo, J., Corchado, J.M.: Agreement technologies and their use in cloud computing environments. *Progress in Artificial Intelligence* 1(4), 277–290 (2012)
5. De la Prieta, F., Rodríguez, S., Bajo, J., Corchado, J.M.: A multiagent system for resource distribution into a Cloud Computing environment. In: Demazeau, Y., Ishida, T., Corchado, J.M., Bajo, J. (eds.) PAAMS 2013. LNCS, vol. 7879, pp. 37–48. Springer, Heidelberg (2013)