On Cloud Deployment of Digital Preservation Environments

Daniel Pop, Marian Neagul, Dana Petcu
Computer Science Department, West University of Timişoara
Bd. Vasile Parvan 4, Timişoara, România
{danielpop,marian,petcu}@info.uvt.ro

ABSTRACT

Although migrating library applications to Cloud environment is not an easy task, many libraries are interested in using Cloud infrastructure services broadly across their businesses, whether is about a Public, Private or Hybrid Cloud. One of the migration expectations is the scalability of digital preservation architectures in Cloud environments. In this paper, we address the scalability of storage and compute platforms, which combine storage of large datasets and their processing. Particularly, we investigate the concrete aspects of managing nodes of the Cloud infrastructure, service provisioning over the infrastructure and the maintenance of the software stack. After presenting some architectural considerations of a scalable deployment system, we detail the prototype developed within a support platform, namely SCAPE.

Categories and Subject Descriptors

H.3.7 [**Digital Libraries**]: Systems issues; C.1.4 [**Parallel Architectures**]: Distributed architectures

General Terms

Applications

Keywords

Cloud computing, Digital preservation

1. INTRODUCTION

The interest in migrating digital preservation systems towards Cloud infrastructures has been recently raised. Several reports like [4, 12, 3] are mentioning such interest from librarians in using Cloud services broadly across their businesses, either via Public-, Private- or Hybrid- Clouds. Moreover, according to [1], archiving in the cloud is a constantly growing business and thus cannot be neglected by digital librarians. A guiding example is the Report on Digital Preservation and Cloud Services [4] published by the Minnesota Historical Society that provides an overview of the strategic and technical issues surrounding the potential use of a wide range of Cloud services for preservation. In this report the authors recognize that there is a gap between what service levels libraries expect and what current public Clouds deliver.

There are numerous advantages offered by Clouds to enterprises business, such as scalability, redundancy, availability, elasticity, or ubiquitous access that would apply for digital preservations as well. Of particular interest for digital preservation is handling of rich metadata for which the Cloud sounds an attractive solution since it exposes a data model of objects that integrates raw data with its user-defined or system-defined metadata as a single unit. Cloud environments have their shortcomings as well: data lock-in, bit reliability (although fixity check is performed upon storing, there is no feature to periodically repeat it), lack of support for auditing, certification and trust that is crucial to businesses that may require regulatory compliance, to name a few. Some of the shortcomings may be alleviated by the promises of Hybrid Cloud [8] that would enable institutions with strict legal regulations on sensitive information handling to store this data in a Private Cloud. Regarding the cost of cloud storage, this is an open debate, some authors asserting that "cloud storage is not even close to cost-competitive with local disk storage for long term preservation purposes" [11], while others are stating exactly the opposite "cloud-based storage is 74% less expensive than in-house" [10].

In [9, 14] the authors emphasise the importance of storage and compute synergy: computational support is needed for preservation, as storage is active over time. In particular the authors of [9] underline that data migration and transformation are an integral part of the preservation digital asset lifecycle, and should be configurable and operable by the client.

However, it is not enough to ensure bit integrity for long term, it must also support the visualization of the content in years to come, i.e. logical preservation. To achieve this aim the digital objects need to be processed (migrated, checked for logical integrity etc.), not only stored in digital media. The development of an open source platform that orchestrates semi-automated workflows for large-scale, heterogeneous collections of complex digital objects is on-going in the framework of the SCAPE[1] project. The SCAlable Preservation Environment aims to provide tools and services for the efficient planning and application of preservation strategies for large-scale, heterogeneous collections of complex digital objects [5]. In particular, the SCAPE platform, earlier described in [14, 13], combines scalable services for planning and execution of institutional preservation strategies that can benefit of deployment on Hybrid Clouds.

In this paper we are addressing the problem of scalability of digital preservation architectures based on Cloud services. We are not focusing here on storage in the Cloud as this has been already addressed in other papers [11, 9], rather we address the issue of scalability of storage and compute platforms that combine storage of large datasets and their processing. Particularly, we are investigating the concrete aspects of adding or removing nodes to/from the Cloud infrastructure, provisioning applications over the infrastructure, and maintenance of the software stack in the context of data or HPC centre operation that raises its specific challenges (sharing of resources, dynamic allocation).

After providing the possible approaches in Section 2, Section 3 presents our architectural considerations for the scalable deployment system. Section 4 details the prototype developed within the SCAPE platform, while few conclusions are outlined in last section of the paper.

2. DEPLOYMENT STRATEGIES

Digital preservation environments require the setup of complex systems comprising multiple software components that need to seamlessly work together. This fact leads to the need for a strategy able to deal with context dependencies in a distributed and dynamically scaling environment on demand without requiring to perform expensive, repetitive operations over again. Hence, we need a consistent packaging and deploying model to manage and sustain the digital preservation components in use in a digital library institution. Packaging and virtualization provide important concepts and tools for the deployment and operation of a scalable infrastructure and are at the core of different possible deployment strategies that are briefly introduced below.

One of the most common and easy to implement strategy is based on pre-configured virtual machine image. It allows customization of the deployment according to each institution's toolset in use. It is then easy to install it on desktop computer or to deploy it in a cluster of multiple machines. In [14] the authors use this approach to provide a basic environment for SCAPE platform, helping its users to experiment with the basic framework and tools built in the context of the project. As described in [9], virtual machine images are used to deploy so called *virtual appliances* on compute clouds (e.g. Amazon EC2, OpenStack Nova) in order to build a PDS Cloud[2]. Although the strategy works fine for small to large deployments, it does not offer any support for further operation and maintenance of the infrastructure. If new or updated packages are required on the infrastructure then a new version of the virtual machine image need to be created and then re-deployed on all machines. This process is overly complicated and error-prone in case of updates of running, 'live' machines.

A typical solution to this problem is based on customized deployment scripts that orchestrate the machine provisioning and the software deployment process on a large number of nodes. Generally this scripts are provider and software specific, requiring multiple versions based on cloud provider, software version and configuration. This type of deployment approach is a tedious and error-prone process, mostly due to testing difficulty and volatility of the operating environment. Nevertheless, it is largely used by organisations as deployment strategy.

More and more organisations, both academic and industrial, with large and complex environments are migrating towards the deployment using specialized *configuration management systems* – such as Puppet[3], Chef[4], CFEngine 3[5] or JuJu[6] among most popular ones – that abstract and automate configuration operations of large scale computer systems. Commonly, configuration management solutions are backward chaining expert systems that aim to bring nodes to a certain state by applying known operations or configuration changes. This strategy provides a more deterministic, reproducible environment suitable for larger deployments that fixes most of the drawbacks of script based deployments, such as scalability and maintenance issues, configuration drift or non-compliance to name just a few. Configuration management systems make it easy for system administrators to provision, configure and manage a heterogeneous infrastructure throughout its lifecycle.

In the next section we present the architectural design and a software toolkit that enables configuration management based deployment of digital preservation environments on hybrid clouds. We were not able to find evidence of similar approaches in the field of preservation environments although the approach is widely used by large companies in domains such as telecommunications, banking and finance or Internet-related services and products.

3. ARCHITECTURAL DESIGN AND DEPLOYMENT TOOLKIT

Figure 1 illustrates the top-level architecture of a digital preservation environment and how the software components are deployed on a Hybrid Cloud infrastructure. The digital preservation environment illustrated here, although based on SCAPE proposal [7], it is generic enough to highlight the complexity of a system that combines storage and processing of digital objects. At the core of the design is the Digital Object Repository (DOR) that stores the digital objects and provides interfaces (API) for all other components of the system. The Execution Platform manages and runs parallelised preservation workflows responsible for expedient and reliable execution of preservation actions on some data set within the DOR, and ensures the validity of the outcome.

The Automated Watch component builds and monitors a

[1]http://www.scape-project.eu
[2]http://www.ensure-fp7.eu
[3]http://www.puppetlabs.com
[4]http://www.getchef.com/chef
[5]http://cfengine.com
[6]https://juju.ubuntu.com

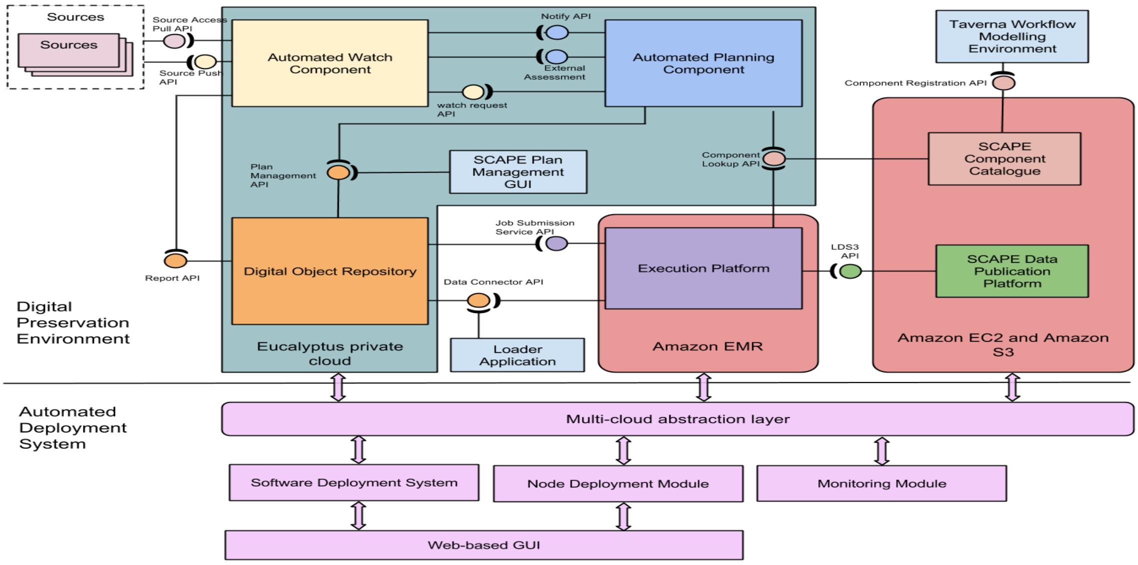

Figure 1: Mapping of digital preservation components on Cloud environments (Loader Application and Taverna Workflow Modelling Environment are run on end-user's computer)

view of the world based on its input sources in relation to institutional policies. The Planning Component is responsible for creating, monitoring and testing preservation plans.

All components and interfaces of this architecture are applications that are bundled in software packages using appropriate package manager tools, such as APT or RPM. The *Software Deployment System* uses customized configuration management solutions in order to orchestrate the packages' installation on target machines. It allows the evolution of digital platform's software packages by providing high-level rule-based recipes describing the tools and relations between them. It enables dynamic allocation of software components to computing resources with minimal human intervention, providing, in this way, a deterministic software deployment process. It ensures that software is deployed as expected by the developers, meeting all required expectations.

Dynamic node allocation in the infrastructure is handled by the *Node Deployment Module* that supports the addition and removal of computing resources on the fly, both on bare-bones hardware and on virtualized resources.

The proposed toolkit supports the deployment on multiple Cloud services provided by different vendors. The heterogeneity allows the user to experiment with diverse technologies and setup, and to examine interoperability across space that will hopefully lead to interoperability across time. To accomplish this objective we introduced a *multi-cloud abstraction layer* service that supports both public and private clouds. It provides a unified API to interact with popular Cloud service providers.

A *Web-based Toolkit GUI* exposes through a friendly user interface the services of Node Deployment Module allowing administrators to easily perform administrative tasks specific to platform deployment on Hybrid Clouds. Moreover, the monitoring of the infrastructure services and quality assurance capabilities allow system administrators to provide a better service to their end-users.

The infrastructure services are offered either by a Public Cloud Provider (CSP), such as Amazon or Rackspace, or by a Private Cloud managed by Eucalyptus or OpenNebula software. For example, Amazon offers compute services (Amazon Elastic Cloud Computing EC2, Amazon Elastic MapReduce EMR), object storage services (Amazon Simple Storage Services S3), block storage services (Amazon Elastic Block Storage), or cheap archival services (Amazon Glacier) that are suited for digital media archives, long-term database backups, and data that must be retained for regulatory compliance.

4. CASE STUDY

In order to evaluate the architectural model proposed in previous section, we implemented a Cloud deployment toolkit in the framework of SCAPE platform, which combines a multitude of technologies and software packages to deliver a scalable environment for digital preservation. The reference implementation of DOR is based on Fedora Commons[7] repositories, while the Execution Platform is composed of Cloudera distribution of Hadoop Apache[8], the Taverna[9] server for executing preservation workflows, plus specific tools and packages for preservation tasks, such as: pagelyzer [6] for Web pages comparison, jpylyzer [15] for validation and feature extraction for JP2 images and xcorrSound [2] for audio comparison. SCOUT[10] is used as Automated Watch component and the preferred choice within the project for Automated Planning component is PLATO[11] planning tool.

[7]http://www.fedora-commons.org
[8]http://hadoop.apache.org
[9]http://www.taverna.org.uk
[10]https://github.com/openplanets/scout
[11]http://www.ifs.tuwien.ac.at/dp/plato

SCAPE Cloud Toolkit[12] (SCT), presented for the first time in this paper, is a prototype that combines state-of-the-art implementations for deployment management modules identified in the architectural design and deploys SCAPE platform components on a Private Cloud based on Eucalyptus[13] software. Eucalyptus has the advantage of supporting Hybrid Clouds compatible with Amazon Web Services (AWS).

We leverage Apache Libcloud[14], an open source cloud interface library, as the basis for the multi-cloud service. It comprises a unified interface (multi-cloud interface component) and a set of drivers that implement the interactions with the individual Cloud infrastructure. The orchestration of software deployment (comprising both installation and configuration) is managed by Puppet configuration management system that employs a Ruby based domain specific language allowing the definition of package deployment recipes in a deterministic and idempotent manner, avoiding occasionally, non-deterministic behaviour of BASH scripts. The idempotent nature of Puppet allows frequent updates of node/package configuration without introducing unexpected changes that might result in system corruption. Coupled with other components like PuppetDB[15], Puppet Dashboard[16], or Facter[17], Puppet offers an effective way for "on the fly" node deployment and cluster wide package reconfiguration.

Currently the toolkit is employed in provisioning components used in SCAPE platform, primarily focusing on SCAPE Execution Platform (Hadoop, Taverna, Fedora Repository and other specific tools) on top of a Private Cloud powered by Eucalyptus software and on Amazon EC2 as Public Cloud provider, with potential extension to other clouds supported by Libcloud.

5. CONCLUSIONS

In this paper we proposed a cloud deployment strategy for digital preservation environments and we presented a prototype of a toolkit that enables it. The toolkit allows administrators to deploy preservation components on Hybrid Clouds in an integrated and portable way, avoiding vendor lock-in by abstracting the provider specific details (autoconfig, networking, etc). The toolkit enables easy extension with new, complex packages and provides a simple interface for users, hiding the unavoidable complexity that might be introduced by cloud environments. It streamlines the deployment of digital preservation platforms by automatizing nodes addition, or tedious tasks such as key-based authentication and it enables reconfiguration, updates and maintenance of installed software components in a way that preserves their integrity and avoids incompatibilities between components.

We are planning to continue with the integration of remaining SCAPE components in order to transform SCAPE platform into a turn-key solution for large-scale digital preservation.

[12]https://bitbucket.org/scapeuvt
[13]http://www.eucalyptus.com
[14]https://libcloud.apache.org
[15]https://docs.puppetlabs.com/puppetdb/latest/
[16]http://projects.puppetlabs.com/projects/dashboard
[17]http://puppetlabs.com/facter

Acknowledgments

The work reported in this paper is partially funded by the European Commission FP7 Programme by grants no. 270137 (SCAPE - Scalable Preservation Environments) and 284595 (HOST - High Performance Computing Service Centre).

6. REFERENCES

[1] E. Anderson. Worldwide storage in the cloud 2011-2015 forecast: The expanding role of public cloud storage services, 2012.

[2] J. Bolette and N. J. Sindahl. Audio quality assurance: An application of cross correlation. In *Proceedings of the 9th International Conference on Preservation of Digital Objects*, pages 144–149. University of Toronto, October 2012.

[3] G. Henry. *Core Infrastructure Considerations for Large Digital Libraries.* Council on Library and Information Resources, Washington, DC 20036, 2012.

[4] Instrumental_Inc. Report on digital preservation and cloud services, 2013.

[5] R. King, R. Schmidt, C. Becker, and S. Schlarb. Scape: Big data meets digital preservation. *ERCIM News*, 2012(89), 2012.

[6] M. T. Law, N. Thome, S. Gançarski, and M. Cord. Structural and visual comparisons for web page archiving. In *Proceedings of the 2012 ACM Symposium on Document Engineering*, pages 117–120, New York, NY, USA, 2012. ACM.

[7] P. May and C. Wilson. Scape D2.2 – Technical Architecture Report v1. `http://www.scape-project.eu/deliverable/d2-2-technical-architecture-report-1`, 2012. [Online; accessed 14-March-2014].

[8] P. Mell and T. Grance. The NIST definition of Cloud Computing, 2011.

[9] S. Rabinovici-Cohen, J. Marberg, K. Nagin, and D. Pease. PDS Cloud: Long term digital preservation in the cloud, 2013.

[10] A. Reichman, R. W. III, and E. Chi. File storage costs less in the cloud than in-house, for infrastructure and operations professionals, 2011.

[11] D. S. H. Rosenthal and D. L. Vargas. Distributed digital preservation in the cloud. *The International Journal of Digital Curation*, 8(1):107–119.

[12] R. Sanchati and G. Kulkarni. Cloud computing in digital and university libraries. *Global Journal of Computer Science and Technology*, 11(12):37–41, 2011.

[13] S. Schlarb. An open source infrastructure for quality assurance and preservation of a large digital book collection. In *Proceedings of the IS&T Archiving*, pages 234–238. IS&T, April 2013.

[14] R. Schmidt. An architectural overview of the SCAPE preservation platform. In *Proceedings of the 9th International Conference on Preservation of Digital Objects*, pages 84–87. University of Toronto, October 2012.

[15] J. van der Knijff, R. van der Ark, and C. Wilson. Improved validation and feature extraction for JP2 (JPEG 2000 Part 1) images: the jpylyzer tool. In *Proceedings of the Archiving 2012*, pages 264–269. Society for Imaging Science and Technology, June 2012.