Render on the cloud: Using Cinelerra on virtualized infrastructures

I. Zablah
Sistema de Difusión de Radio y Televisión
Universidad Nacional Autónoma de Honduras
Tegucigalpa, Honduras
mrzablah@unah.tv

A. Garcia-Loureiro, F. Gomez-Folgar and T.F. Pena
Centro de Investigación en Tecnoloxías da Información (CITIUS)
University of Santiago de Compostela
Santiago de Compostela, Spain
(antonio.garcia.loureiro, fernando.gomez.folgar, tf.pena)@usc.es

***Abstract*—Nowadays, the learning-teaching processes are being supported by the use of new technologies, including email, chat, online conferencing, online activities and videoconferencing. With the revolution of high definition television, video producers require more efficiency in the production and post-production tasks. There are several software available for this purpose including commercial solutions, often very expensive, and open-source ones, such as Cinelerra.**

This paper proposes a cloud infrastructure for using Cinelerra, a community developed version of non-linear video editor, and how educational institutions can use this. The main idea is to reuse its computational power for editing or creating educational videos without the need of acquiring a dedicated hardware infrastructure, employing non dedicated resources, such as the computer labs, or desktop computers to fix the most common time consuming problem: the rendering. The performance of the proposed infrastructure is also presented in this paper.

***Keywords*-Cinelerra; rendering; cloud; broadcast;**

I. Introduction

Recently, technologies [1] such as email, chat, online conferencing, online activities and videoconferencing were incorporated to support teaching and learning processes. Today we live the biggest revolution on computing, multimedia and TV, since the invention of color broadcast in the beginning of fifties decade [2]. Currently, teachers and students can watch videos on a variety of mediums, from mobile phones, computers and high definitions screens. These changes goes on the hand with the evolution of the information technology. They used the computing power of the new computers to create videos of superior quality and complexity from a bunch of source feeds.

It is well-known rendering is a time consuming task. Usually, the companies and training centers use expensive resources to reduce the edition time necessary to prepare the videos. In an educational environment, it could be better to create high quality videos without the need of incurring in license cost and without acquiring expensive dedicated hardware resources. This paper focus on how it is possible to reuse the existing computer hardware of educational institutions, such as schools, colleges and faculties to create educational videos, by implementing cloud and virtualization technology [3].

The cloud technology can be defined as some kind of parallel and distributed system [4], conformed by a lot of interconnected Virtual Machines (VM), or guest systems, providing dynamically computational resources on demand as a unified one. They are based on a Service Level Agreement (SLA) [5]. This technology is growing very fast as well as the computer resources that support it, specially the Service Oriented Architecture (SOA) [6] and the virtualization technology [3], using both hardware and software resources. The virtualization technology is the cornerstone of the cloud, as well known as Infrastructure as a Service (IaaS) [7]. The cloud technology can be used with different objectives. In our case we are interested to know the advantages of use it on video rendering process using the Cinelerra [8] application on a virtualized environment provided by a hypervisor layer.

This paper is organized as follows. The section II describes the implementation of the proposed infrastructure used to install Cinelerra and how use it to improve the rendering process. The section III describes two cases of study for the proposed infrastructure. The Section IV includes the performance evaluation of the proposed infrastructure performing several rendering tests. Finally, the conclusions of this paper are drawn in section V.

II. Proposed Implementation

The proposed infrastructure for executing Cinelerra, shown in Fig. 1, is composed by several types of components: physical compute nodes, a virtual master node, virtual render nodes and a Network File System (NFS), shared by the virtual cluster.

The physical compute nodes are an Intel Core I7-2820QM processor with a clock speed of 2.30 GHz and 8 GB of DDR3 RAM, with the VirtualBox [9] 4.1.14 hypervisor installed. This processor reports eight cores (with hyperthreading enabled). The master and the render nodes are VirtualBox VMs with two cores per VM, with 1 GB of RAM, taking advantage of the hardware virtualization extensions. The virtual cluster infrastructure (virtual master and virtual compute nodes) uses Ubuntu 12.04 LTS 64 bits as guest operating system, employing a Gigabit Ethernet interface. The NFS stores input and output video resources.

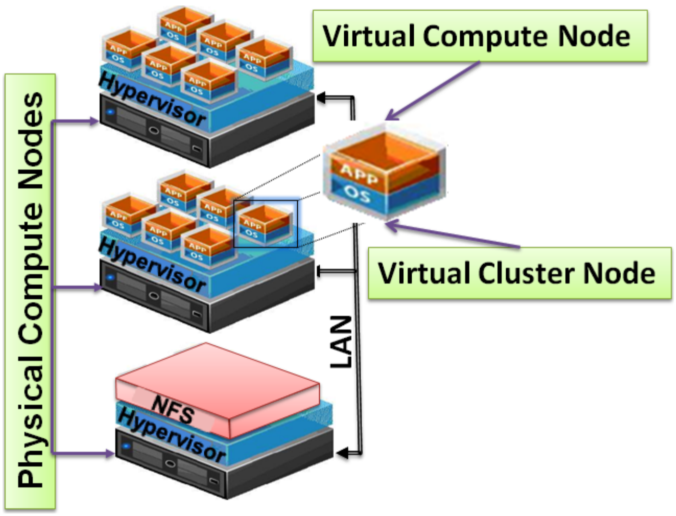

Figure 1. Schematic view of the implementation.

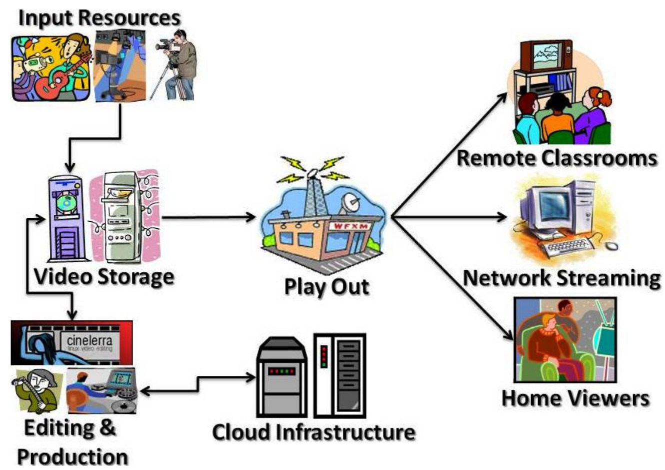

Figure 2. Use of Cinelerra at UNAH TV station.

III. Cases of Study

Two cases of study were proposed as examples of use of the infrastructure described previously. The first one is the implementation of Cinelerra over the infrastructure available at the National Autonomous University of Honduras (UNAH). The second one is using Cinelerra in a training center.

A. *Cinelerra in the public UNAH TV channel*

The infrastructure proposed previously can be implemented at UNAH in order to help the process of producing television content for broadcast on the public TV channel owned by the university, reducing the required time to obtain the final content for airing. This infrastructure takes advantage of the virtualization technology and the Infrastructure as a Service (IaaS) paradigm, allowing reusing the hardware available at UNAH for multiple applications.

Fig. 2 shows the work-flow that must be follow with the main objective to incorporate the advantages of rendering on the cloud. As shown in the figure, at UNAH TV station all the video and TV production starts with the introduction of the media content into the broadcasting system through a common point called *Input Resources* or *Ingest*. Here employing a serie of procedures, videos from a camera, studio, satellite, DVD, tape, etc. are converted in computer video files accompanied with a standardized metadata files that describes their content and properties.

As a result of the previous process, in first place, the video files are stored in the *Video Storage*, that is composed by a group of NFS servers configured with redundant arrays of disk drives. In second place, the metadata is stored in a database server. The *Video Storage* can be organized in hierarchical levels to distinguish videos of different origin, source, resources, media support and the edited ready to air videos.

In the *Editing & Production* department, the editors, equipped with powerful workstations, access to the hierarchical Video Storage to take the required resources that will be incorporated into the Cinelerra video project, including sounds, videos, images, and production scenarios. Once the video is edited, is ready for starting the rendering stage employing the computational power available at the *Cloud infrastructure* that is composed by non-dedicated hardware resources. The rendering stage ends with the creation of a final version of the ready to air TV program. The results of the rendering process are stored in the *Video Storage*, which can be taken by the operators of the *Play Out* department to be transmitted and make them available to final users. Those contents can be also accessed by the academic staff from the classroom using the existing university network, or to the general public employing Internet streaming.

B. *Cinelerra as a tool in educational environments*

This scenario proposed the use of Cinelerra as a tool in educational environments employing a cloud composed by non dedicated hardware resources. In this case, the cloud is employed as a rendering queue where the projects will be processed like in a batch system to get the ready to air videos. This case of study is shown in Fig. 3. As we can see, students create non-linear edition projects using Cinelerra in their workstations. Source videos, transitions effects, and the additional necessary media compose those projects. When the edition process is finished, the project is stored in the Master Node. This element shares the media directory employing the Network File System (NFS) protocol and has also Cinelerra installed as a render queue that manages the jobs pending of being rendered. These jobs are dispatched to the Virtual Render Nodes in which the ready to air video is created. The Virtual Render Nodes are Cinelerra enabled VMs that mount the NFS directories shared by the Master Node. The virtual Render Nodes can be executed in the

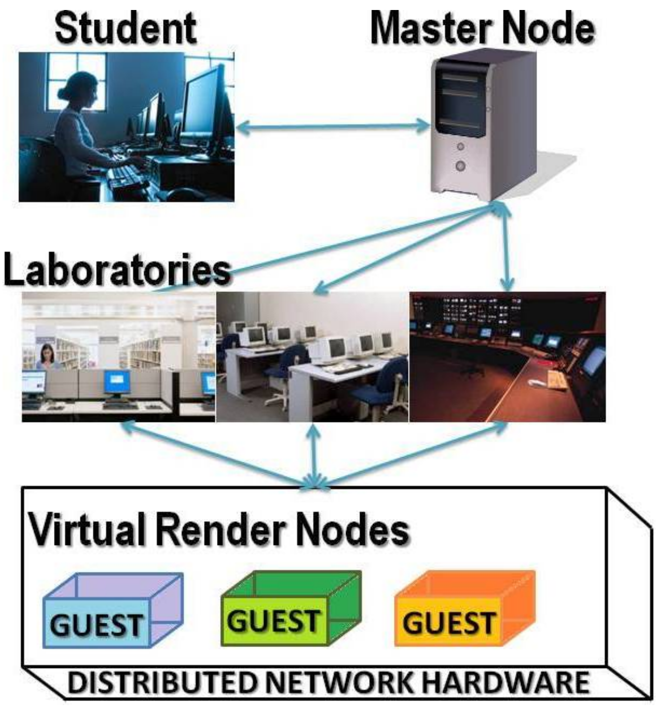

Figure 3. Use of Cinelerra by students with virtual render nodes located on remote computers on other rooms.

computer laboratories of the educational institutions. The video created by the rendering process is finally stored in the NFS shared directory of the Master Node. Notice that the student's workstations, the Master Node and the Virtual Render Nodes must be accessible and interconnected by a communications network.

The main goal of this scenario is the possibility to reuse the idle computational power available in the computer labs elastically, as they are available, to reduce the necessary time to get the rendering process finished.

IV. Performance Evaluation

In this section we included the performance analysis of Cinelerra over the proposed infrastructure, as described previously.

Cinelerra was used for editing and rendering video on resolutions of 720p and 1080i, with input/output streams used for high definition (HD) television. As a common standard for audio/video the codec MPEG4 [10] was employed. It has been on video editing world from last years of 90s decade.

To evaluate if the proposed cloud infrastructure could be considered as a good option for rendering, we prepare some tests. The first one was an output video of 40 seconds length with 1080/60p resolution. The second one was a video of 30 minutes at 720p. The third one was a video of 30 minutes at 720p, composed by thirty feeds of one minute without transitions or effects. We run the render tests only in the master with two cores, later with one and two nodes,

Figure 4. Performance with 40 seconds of video at 1080/60p.

employing two cores each one. The purpose of the tests is to know how much time is needed in every configuration and how the render process could be beneficed or penalized.

When several nodes are used, we set the option in Cinelerra that allows to automatically dividing the job in several parts. Cinelerra itself splits the jobs trying to give the same number of render parts to all machines (including the master). In this way, the number of files generated scale with the number of nodes employed (from one when only the master rendering, two, four, six, eight up to sixteen). This method avoids that some rendering nodes were unused.

For the first test, the video of 40 s at 1080/60p, we can see the obtained results in Fig. 4. The master needed more time to finish than the other configurations, except when we used the master with one node and sixteen output files that required 287 s, the worst performance in this test. The best time was obtained employing the master and two nodes using eight files, requiring 125 s only. Employing the same combination with four files the time required was 134 s. Using the master and one node, the best result was quite the same using two and four output files; this was 153 s, this is 28 s more than the best result of all the present test.

Our second test was a project with an output video of 1800 s (30 minutes) at a standard resolution of 720p. In this project we joined some feeds and added transitions between every resource, applying basic effects. The rendering process took 6487 s for the master. This was the worst result on the present project. After adding one node to help the master, the efficiency was improved. The best result in this combination was 3605 s, allowing Cinelerra dividing the output render job in six files. After adding a second node, the shortest and best time was obtained employing eight output files that took 3072 s only. The results are depicted in Fig. 5. Because the previous video project are not homogeneous and contains

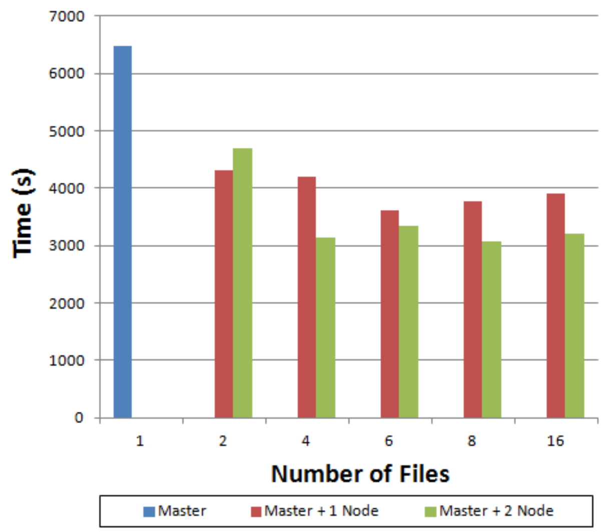

Figure 5. Performance with 30 minutes of video at 720p.

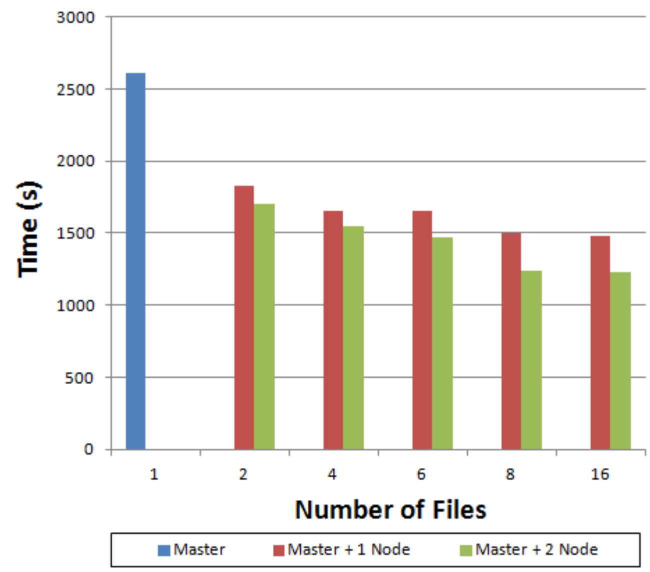

Figure 6. Performance with 30 minutes of video at 1080/60p. Shortest bars, better performance.

several types of effects and transitions, which mainly differ in complexity, we consider to be very interesting to know how the infrastructure proposed would behave to complete a project in which resources were homogeneous, allowing the nodes performing similar computational tasks. The third test, as shown in Fig. 6, was prepared employing 30 feeds one minute of video at 1080i without employing transitions. In this case, the better time was obtained with the master and two nodes, employing sixteen output files that required 1226 s The result is very similar employing eight files. The render using the master required 2614 s, the worst elapsed time in the present measures. Working with the master and one node, the best time was 1477 s, using with sixteen files. We got similar results when this configuration is used with eight output files as the difference was 24 s only.

V. Conclusion

The use of cloud technologies to create high quality videos in an educational environment is feasible without the need of incurring in license cost and without the need of acquire expensive dedicated hardware resources. This paper analysed a cloud infrastructure for using Cinelerra, a community developed version of non-linear video editor. As shown, this software could be used and employed by educational institutions for teaching video techniques and to create educational material. The main goal is to use, in a more efficiently way, the computational resources managing them as a virtual cloud infrastructure for rendering purposes. This avoids the need of acquiring a rendering dedicated hardware infrastructure allowing to reuse existing computer labs or desktop computers making them available to the students where they can easily render their videos.

The infrastructure and use cases proposed provide better performance when working with large projects. Therefore, if we want to render very small videos, the computing capacity available on workstation used by the students or editors will be sufficient to successfully complete the jobs.

It is important to mention that the proposed solution, based on cloud rendering, will be helpful in the process of creating multimedia for the TV station of the National Autonomous University of Honduras, Therefore, we can provide to end users, in a short period, large amount of quality television videos and multimedia using the idle computing power available.

The most important aspect of improving the rendering process is to have a greater chance to train better technical experts and to create fastest multimedia content, therefore, the students will have more computational resources to develop their projects and ideas, but without incurring in the cost of acquiring a dedicated infrastructure.

Acknowledgment

Part of this job was funded through the Fiduciary Fund of the Japan's government - UNESCO with the program Keizo Obuchi 2011-2012, and by FEDER funds and Xunta de Galicia under project 09TIC001CT, contract 2010/28, and by Spanish Government (MCYT) under project TEC2010-17320.

References

[1] Schar, S.G. and Krueger, H.; *Using new learning technologies with multimedia.* IEEE Multimedia, Number 3, Volume 7, Pages 40-51, 2000. DOI:10.1109/93.879767, http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=879767l

[2] WIRED: The First Color TV; *March 25, 1954: RCA TVs Get the Color for Money*. http://www.wired.com/science/discoveries/news/2008/03/dayintech_0325

[3] Bhardwaj, Sushil, Leena Jain, Sandeep Jainl; *CLOUD COMPUTING: A STUDY OF INFRASTRUCTURE AS A SERVICE (IAAS)*. International Journal of Engineering and Information Technology, Pages 60-64, 2010.

[4] J. J. Rehr, J. P. Gardner, M. Prange, L. Svec, and F. Vila. *Scientific Computing in the Cloud*, December - 2008. http://arxiv.org/abs/0901.0029

[5] Buyya, Rajkumar and Garg, Saurabh Kumar and Calheiros, Rodrigo N.; *SLA-Oriented Resource Provisioning for Cloud Computing: Challenges, Architecture, and Solutions*. 2012. arXiv:1201.4522, http://arxiv.org/abs/1201.4522

[6] Hassan, Qusay F.; *Aspects of SOA: An Entry Point for Starters*. February, 2012. arXiv:1202.6623, http://arxiv.org/abs/1202.6623

[7] Ambrust, Michael et al.; *Above the Clouds: A Berkeley View of Cloud Computing*. EECS Department, University of California, 2009. http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html

[8] Cinelerra-CV: Video editing for Linux. http://cinelerra.org/

[9] VirtualBox. https://www.virtualbox.org/

[10] MPEG-4 - The Media Standard. 2002. http://www.m4if.org/public/documents/vault/m4-out-20027.pdf