

An Architecture for Public and Open Submission Systems in the Cloud

Marcello Azambuja¹, Rafael Pereira¹, Karin Breitman², Markus Endler²

¹WebMedia – Globo.com

Av. das Américas 700 Bloco 2A, 1º andar – 22640-100 – Rio de Janeiro, RJ – Brazil

²Departamento de Informática – PUC-Rio

Rua Marquês de São Vicente, 225 – 22453-900 – Rio de Janeiro, RJ – Brazil

azambuja@ieee.org, rafael.pereira@corp.globo.com,
{karin,endler}@inf.puc-rio.br

Abstract. *The advent of the Internet poses great challenges to the design of public submission systems as it eliminates traditional barriers, such as geographical location and costs associated to physical media and mailing. With open global access, it is very hard to estimate storage space and processing power required by this class of applications. In this paper we argue in favor of a Cloud Computing solution, and propose a general architecture in which to built open access, public submission systems. Furthermore, we demonstrate the feasibility of our approach with a real video submission application, where candidates that want to take part in a nationwide reality show can register by submitting a personal video clip.*

1. Problems Addressed

The Brazilian Big Brother reality show is broadcasted by open TV network with an audience of more than eighty million people simultaneously. The idea behind this reality show is to portray the life of 16 random anonymous people while living together under the same roof, for a total period of three months. They are isolated from the outside world but are continuously monitored by television cameras. The housemates try to win a cash prize by avoiding periodic evictions from the house.

With technological advances the application process evolved from sending a video tape by postal mail to uploading a digital video using the Internet. Due to legal reasons, videos can not be hosted in websites such as YouTube or Vimeo. Applicants are allowed to send videos in the video format of their choice. These must be stored until the end of the selection process (three months). All the videos need to be transcoded to a standard format, so that the TV show's production team is spared from the hassle of having to deal with a plethora of video formats and different codecs.

The system should be able to receive a very large number of videos during the three-month application process. With the new digital process it is expected that more than 100,000 videos; about 60% of the total submission is uploaded during the last week before the deadline.

2. Relevance to SBRC

Investment in infrastructure for high peak situations for a short period of time are usually a waste of money and resources as most of the time the resources will not be used. In what follows we will argue that Cloud Computing [Armbrust, et al. 2009] technology provides the necessary requirements in which to provide a viable solution. It provides the necessary infrastructure in which to develop submission applications in which both storage and processing needs can be dimensioned as needed. For this purpose we propose a general architecture for open, public submission systems – thus allowing up and down scalability rapidly responding to external factors.

The relevance to SBRC of this demo is the use of Cloud Computing [Vogels, 2008] to solve a real world problem with a general purpose architecture that could be re-used in different situations. This architecture provides the necessary flexibility to be used in a wide range of applications [Miller, 2008] that deals with large dataset processing, such as text corpus processing, audio recognition, and, as described in this demo, for mass video transcoding, and, when deployed in a Cloud platform, providing a dynamic and efficient resource usage, which might be a critical factor to business success.

3. Uniqueness of Design and Implementation

In this section, we describe the proposed architecture for large, user generated content, file submission and processing systems using Cloud Computing. A few specific characteristics leverage the use of a Cloud Computer architecture for this project in particular:

- Uncertainty in how much storage and processing capacity will be needed;
- Resources will be needed during the application and selection processes only. After this short period, all storage and computing resources would be idle;
- Few but extreme high peak situations where the infrastructure will need to scale up – 60% of the videos are expected to be sent in the last week;

The proposed architecture uses the cloud to store and process all this content, and to provide storage availability and scale resources as needed. All user content is received through a website where video files can be uploaded without restriction regarding the file extension, or video format/codec.

The demo is based on Amazon's Cloud Computing platform and we make use of the following services:

Amazon S3 – Amazon Simple Storage Service is cloud-based persistent storage and operates independently from other Amazon services. It can be used to upload data in the cloud and pull it back out.

Amazon EC2 – Amazon Elastic Compute Cloud is a web service [Zhang, 2007] that provides resizable compute capacity in the cloud. It provides an API for provisioning, managing, and deprovisioning virtual servers inside the Amazon cloud. It's the heart of the cloud and allows remote deployment of virtual server with a single web service call.

Amazon SQS – Amazon Simple Queue Service is a highly scalable, reliable, hosted queue for storing messages as they travel between computers. It can be used to move data between distributed components of an application that perform different tasks, without losing messages or requiring the components to be always available.

As each submission is made, the video file is stored in Amazon S3 and a message is written in SQS Queue with relevant information so that proper processing of the job can be done. An EC2 instance is created to process the new submission using the information contained in the SQS queue. The message contains relevant information for an EC2 instance process a new job, consisting of transcoding the user's video to a standard format, bitrate and specific codec (MPEG4/h.264/AAC). The output video should be easily reproduced by any video player, e.g., Adobe's Flash video player.

Figure 1 illustrates the complete workflow of the system.

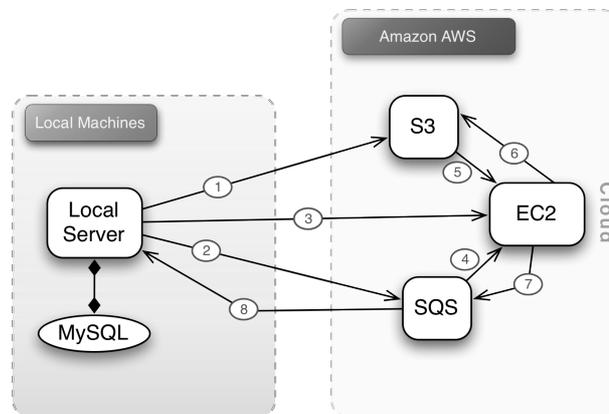


Figure 1. Architecture for the public submission system.

We detail the process in a following basic steps:

1. Video submitted by user is stored in Amazon S3;
2. Local server writes the message in the input queue of SQS detailing the job to be done;
3. Local server creates a new EC2 instance to process the job;
4. EC2 instance reads the message from the input queue;
5. Based on the data of the message the input video is retrieved from S3 and stored locally in the EC2 instance;
6. Video is transcoded by EC2 and the generated output is stored in S3;
7. EC2 instance writes a message in the output queue describing the work performed;
8. Confirmation of the work completed is read by the local server from SQS output queue.

The local server illustrated in the picture is the web application responsible for receiving the user generated content.

Messages use the basic structure format used by mail messages and HTTP headers defined in RFC-822 [Crocker, 1982]. Input messages are as follows:

```

Bucket: bbb.video
InputKey: f84e4a21b571abc69baf2277d193e596
Date: Tue, 29 Oct 2009 17:21:21 GMT
OriginalFileName: EF_512.AVI
Size: 25203791

```

Where *Bucket* and *InputKey* are the identifiers of the file in the S3 infrastructure, and *OriginalFileName* is the source filename. Webservices were implemented using the Boto library (see References).

The output message is defined as:

```

Bucket: bbb.video
InputKey: f84e4a21b571abc69baf2277d193e596
Date: Tue, 29 Oct 2009 17:21:21 GMT
OriginalFileName: EF_512.AVI
Size: 25203791
OutputKey: e69e376be5af6f88f81d3e31adf27988;type=video/quicktime
Host: ec2-75-101-237-173
Service-Read: Tue, 29 Oct 2009 01:28:14 GMT
Service-Write: Tue, 29 Oct 2009 01:28:27 GMT

```

Where we also define the hostname of the EC2 instance that processed the job and the timestamps when the job was received and when it finished.

The EC2 instance that is launched uses a specific Amazon Machine Image (AMI) created with all the dependencies necessary to process the video. That includes an updated version of the Linux kernel, git to retrieve the latest source code available for this framework, Python and the FFmpeg software which does the actual video transcoding.

Once the AMI image is instantiated it reads a configuration file that keeps parameters as:

- command line and arguments – in this case the ffmpeg command;
- maximum processing time before marking the job as dead;
- input queue name to read SQS messages;
- output queue name to store SQS messages;
- maximum number of retries in case of error;
- notification e-mail (for debugging purposes);
- python class to be invoked for the processing.

Due to the generalization of the configuration file, the framework can be setup to a variety of other purposes not restricted to video transcoding.

4. Underlying Implementation Techniques and Used Technologies

In the complete system there are three different sub-systems:

- Web application for receiving video files from users in the Internet;
- Back-end system to manage the cloud infrastructure, creating EC2 instances, writing/reading in SQS and storing files in S3;
- Video transcoding application of the received content – run in the cloud instance;

The web application was developed using PHP and is the only system with an interface to the end user – the website itself.

The back-end system was written in Python using Boto library (see Reference) to consume Amazon's web services to manage the cloud infrastructure. The use of Amazon's cloud platform allowed the architecture to be scalable and elastic to fulfill the high peak demand as needed.

Finally, for the transcoding of the video we used FFmpeg, which is a complete solution to convert audio and video. FFmpeg's libavcodec provides support to a great number of different video formats and codecs.

5. Description of Presentation

To begin the submission process the user needs to create an account in the reality show's (Big Brother Brasil) website.



Figure 2. Big Brother Brasil's web page.

The user can either use his existing account or create a new. The user account is a requirement so that we can guarantee that the rights of the content are not being violated by the end user – the user must accept an agreement claiming he is the owner of the content.

The screenshot shows a web browser window with the URL `http://login.dev.globo.com/login/1129`. The page header includes the Globo.com logo and a navigation bar. The main content area is divided into two columns. The left column features a 'Big Brother Brasil' banner with the text 'As inscrições estão abertas. Fique de olho!' and a 'User Generated Content' section with a 'Send us your video!' link. The right column contains a 'cadastroglobo.com' sidebar with a login form (Email/Login: 'azambuja', Password: [empty]), a 'Keep me signed in' checkbox, a 'sign in' button, and a 'forgot password' link. Below the login form is a 'Sign up!' section with the text 'cadastroglobo.com?' and a 'cadastre-se agora »' button.

Figure 3. Sign in / sign up web page.

Some meta-data information needs to be filled and the video chosen from his local filesystem.

The screenshot shows a 'User Generated Content' submission form. The form is divided into several sections:

- content**: Includes fields for 'subject: Bola Cheia, Bola Murcha', 'title: Title of the Video', 'keywords: keywords video ugc', and 'summary: Summary and description of the video.' There is an 'edit' link next to the section header.
- personal information**: Includes fields for 'name: Marcello L. Azambuja', 'email: azambuja@gmail.com', and 'phone: 55 21 8153-9795'. There is an 'edit' link next to the section header.
- video**: Includes a 'video*' field with a list of supported formats (3g2, 3gp, 3gp2, 3gpp, asf, avi, divx, dv, dvx, f4v, h263, m4e, m4v, wmv, mov, movie, mp4, mpg, mpeg, qt and rm). Below this is a file selection field with the path '/Users/azambuja/videos/fo' and a 'Browse...' button.
- description***: A text input field with a character count of '200'.
- confirm submission**: Includes a text input field with the instruction 'Fill the field below with the text from the image.*' and a small image of the word 'imbuzada'.

 At the bottom of the form is a 'submit' button. A note at the bottom left states '*campos obrigatórios'.

Figure 4. Submission form web page.

Once the video is received by the system, the Cloud is taken into place. Figure 5 shows the instances created in Amazon EC2 as the videos were being received.

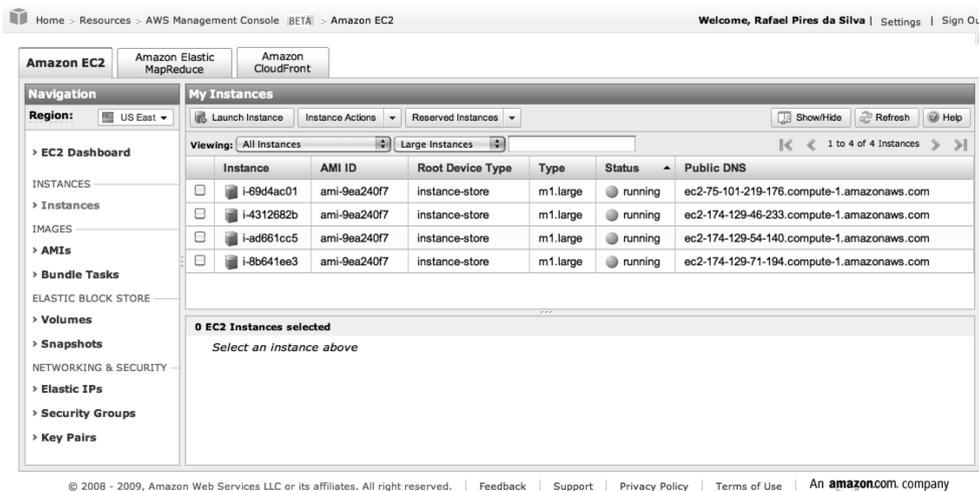


Figure 5. Amazon AWS Management Console.

Once the transcoding process is completed the instances are shutdown to avoid wasting computing resources from the EC2.

If we were to calculate how much money would be spent to process 100,000 videos with an average size of 15MB, using the small EC2 instance we can process a video in 50% real time, that would require 834 hours of CPU running.

Table 1. Cost analysis of transcoding solution in the Cloud for 100,000 videos.

Storage	1.5 TB	US\$0.15 / GB	US\$ 225.00
Transfer	1.5 GB	US\$0.14 / GB	US\$ 210.00
Messages	200,000	US\$0.01 / 10,000 requests	US\$ 0.20
Computer Resources	834 hours	US\$0.085 / hour	US\$ 70.90
Total			US\$ 506.10

A total of US\$506.10 for transcoding and storing 100,000 videos – that's not even a penny for each video.

Adding to that the fact that no up-front investment and deployment of infrastructure was needed, neither a precise estimation on the expected load of the system we can conclude that Cloud Computing is an excellent solution in this specific scenario. It is important to remark that economical viability of the proposed architecture is such that enables it to quickly deploy at great reduction of the TCO, typical of Cloud Computing implementation [Walker, 2009].

6. Hardware and Presentation Requirements

If possible we'd like to use the presenter's laptop to run the demo. In this case we only need a stable network connection so that we can connect to Amazon Cloud platform. To

access Amazon we need to be able to connect to ports TCP/22 and TCP/80 for the following IP ranges:

- * 216.182.224.0/20
- * 72.44.32.0/19
- * 67.202.0.0/18
- * 75.101.128.0/17
- * 174.129.0.0/16
- * 204.236.224.0/19

7. A Note on Code Availability

Because the implementation of this tool relies on computational resources located on a public Cloud infrastructure, as opposed to stand-alone applications, it is not possible to provide a single URL where the tool is made available.

In the hopes that the following provides as a suitable alternative to the referees, we invite all to check:

1. the server image used in the demonstration described in section 5. Please note that one must be registered in the AWS Console / AMI to be able to instantiate a machine.

http://sbrc-cloud-video.s3.amazonaws.com/submission_framework.manifest.xml

2. a video of the Submission System framework demo:

http://sbrc-cloud-video.s3.amazonaws.com/demo_cloud_submission.mp4

3. other forums where our related research is presented:

- a. CloudSlam'10: Architectures for Distributed High Performance Video Processing in the Cloud –

<http://cloudslam10.com/taxonomy/term/33?page=4>

- b. Microsoft Cloud Futures Conference: Cloud TV –

<http://research.microsoft.com/en-us/events/cloudfutures2010>

- c. Cloud Computing Brazil: Video Processing in the Cloud –

<http://www.ccbrazil.com.br>

- d. Conseg 2010: Desafios e oportunidades em Computação na Nuvem: do Big Brother ao Voto Eletrônico – III Congresso Internacional de Software Livre e Governo Eletrônico –

<http://www.consegi.gov.br>

References

Armbrust, M., Fox, M., Griffith, R., et al. (2009) “Above the Clouds: A Berkeley View of Cloud Computing”, In: University of California at Berkeley Technical Report no. UCB/EECS-2009-28, pp. 6-7, February 10, 2009.

- Boto Library, Python interface to Amazon Web Services – <http://code.google.com/p/boto/>
- Cha, M., Kwak, H., Rodriguez, P., Ahn, Y., Moon, S., (2007) “I Tube, You Tube, Everybody Tubes: Analyzing The World’s Largest User Generated Content Video System”, Internet Measurement Conference, ACM, 2007.
- Converting video formats with FFmpeg, Linux Journal archive – Issue 146, June 2006, pp. 10
- Walker, E. “The Real Cost of a CPU Hour”, IEEE Computer, vol. 42, no. 4, pp. 35-41, Apr. 2009, doi:10.1109/MC.2009.135
- Information technology – Coding of audio-visual objects – Part 10: Advanced Video Coding, ISO/IEC 14496-10:2003 – http://www.iso.org/iso/iso_catalogue/catalogue_ics/catalogue_detail_ics.html?csnumber=37729
- Jensen, C. S., Vicente, C. R., Wind, R. (2008) “User-Generated Content: The Case for Mobile Services”, In: IEEE Computer, vol. 41, no. 12, pp. 116-119, Dec., 2008.
- Miller, M. (2008) “Cloud Computing: Web-Based Applications That Change the Way You Work and Collaborate Online”, Que, Aug. 2008.
- Crocker, D. (1982) ”Standard for the format of ARPA Internet text messages” – Request For Comments 822 – <http://www.ietf.org/rfc/rfc0822.txt>
- Vogels, W. (2008) “A Head in the Clouds – The Power of Infrastructure as a Service”, In: First workshop on Cloud Computing in Applications (CCA '08), Oct., 2008.
- Zhang, L., (2004) RESTful Web Services. Web Services, Architecture Seminar, University of Helsinki, Department of Computer Science, 2004.