

CLOUD MY TASK – A PEER-TO-PEER DISTRIBUTED PYTHON SCRIPT EXECUTION SERVICE

*Eng. Daniel RIZEA¹, Eng. Daniela ENE², Eng. Rafaela VOICULESCU³,
Lecturer Eng. Ph.D. Mugurel Ionut ANDREICA⁴*

Abstract

In this paper we present „Cloud My Task”, a distributed Python script execution service. The service can receive Python scripts from the clients and execute them on its instances, according to their load. The service is fault-tolerant : the failure of any single instance will not disrupt the service availability. The instances are connected with each other in a peer-to-peer tree-like topology. Experimental results show that the service scales well and is able to balance the load across the available service instances.

Keywords: *Cloud, task scheduling, distributed service, load balancing, P2P, topology.*

1. INTRODUCTION

In this paper we introduce „Cloud My Task”, a distributed Python script execution service implemented in the Java programming language. The service can distribute CPU-intensive Python scripts in a cloud environment balancing the computations and returning the results to clients as soon as they are available.

The motivation for focusing on Python scripts/programs is that Python is an easy and intuitive language that can be learnt very fast so that non-programmers can develop scripts in a matter of minutes. Scripts in general can run for long periods of time and may need CPU resources which are not available locally. Running these scripts in a cloud environment like „Cloud My Task” will alleviate the need for available local resources. „Cloud My Task” will also store the results of running the script for a period of time, so that these results may be retrieved by the users later.

The rest of this paper is structured as follows. In Section 2 we present the main features of the „Cloud My Task” service. In Section 3 we discuss service implementation details. In Section 4 we present the architectural design of the service. In Section 5 we discuss the TPSI topology. In Section 6 we present experimental results. In Section 7 we discuss related work and in Section 8 we conclude and discuss future work.

2. FEATURES OF THE „CLOUD MY TASK” SERVICE

The „Cloud My Task” service can run CPU-intensive Python scripts in a cloud environment with a high uptime. The machines composing the service are geographically close, usually in the same cluster or local area network. This way, both point-to-point and IP multicast communication between machines can be employed. The service instances are interconnected into a peer-to-peer topology, where each instance has a limited number of neighbors. Moreover, the service is fault-tolerant - the failure of a machine will not lead to the disruption of the service.

In order to use the „Cloud My Task” service, a client must communicate with an instance *IA* of the service (running on one of the machines). We will not address the service discovery

¹ Politehnica University of Bucharest, Computer Science Department, e-mail: danielrizea27@gmail.com

² Politehnica University of Bucharest, Computer Science Department, e-mail: daniela.a.ene@gmail.com

³ Politehnica University of Bucharest, Computer Science Department, e-mail: rafaela.voiculescu@gmail.com

⁴ Politehnica University of Bucharest, Computer Science Department, e-mail: mugurelionut@gmail.com

problem in this paper (i.e. we will not focus on mechanisms which allow the client to find an instance of the service to communicate with). The client sends its Python script to the service instance in order to be run. If *IA* is not too loaded, it will add the script to its processing queue and the script will be run as soon as possible. If the load on the instance *IA* is too high and the instance cannot process the demand right away, it will try to find another instance of the service with a load below the critical threshold. If such an instance *IB* is found, then the task will be forwarded to that instance and when the task will be done the result will be returned to *IA*. This way, the client only ever needs to communicate with *IA*, although its script may be run on a different service instance. At a certain point when all the instances are highly loaded, a script will wait for the CPUs to be free (in a special queue). The instance *IA* will sense when this happens and will pass the waiting script to a less loaded instance of the service (or run the script itself if it becomes less loaded).

The output for a script will be stored for a certain period of time. The client will get a unique id associated with that script and will be able to query the service for the result of the script.

3. SERVICE IMPLEMENTATION DETAILS

The service implementation will be analysed based on the interaction between the service clients and service instances. We will have client-task processing service instance (TPSI) and TPSI-TPSI interactions. We will first describe the main components of our services and then describe the interactions between them.

3.1. Service Components

3.1.1. Clients

A public interface will be offered to client application in order to be able to use our service. They will mostly use one method:

```
public Integer submitJob(ClientID, QoS type, Script) { ... }
```

The command will return the script job ID in the system. Clients will then be able to query for the status of the script they have submitted using the script job ID obtained from submitting the script, using the method below:

```
public String getScriptJobStatus(ScriptJobID) { ... }
```

This method will return the script result or will return the status of the current script. (Pending, Running, Waiting, etc.).

3.1.2. Task Processing Service Instance (TPSI)

This is the main part of our service. The TPSIs execute Python scripts and returns their results to client application. A TPSI communicates with client applications to get script jobs, it communicates with other TPSI instances to pass tasks so that load balancing is achieved in the system and it also communicates with the Central Supervisory Unit to query and find out if users are in a banned list.

3.1.3. Central Supervisory Unit

This is a supervisory unit that will periodically collect system information and overall state. For now it will hold a client blacklist that will prevent clients from achieving a DoS attack on the distributed system. Each time a TPSI senses the possibility of an attack it will

reject the client's job requests and call *addToBlackListRequest* to the Central Supervisory Unit.

3.2. Interactions between components

We will present the main types of **requests** based on the components that are interacting. The requests are broken down into smaller steps. Each of these steps composes a stage of the request processing pipeline and will be executed by a thread from a pool of threads (either created on demand or statically created in the beginning). The results from a previous stage will be submitted in the job queue of the next stage along with other context information. Based on the **3 possible interactions** we will have **3 main** requests types available.

3.2.1. Client job request

This type of request is submitted from a client to a Task Processing Service Instance (TPSI). This request involves the submission of a Python script along with other identifiers. The processing stages of this type of request are shown below (and also presented in Fig. 1):

1. Receive the Request (get the request and the Python script).
2. Decode the Request (decide the request type and what logical flow to follow).
3. Submit request to Central Supervisory Unit to check if client is banned from the service and wait for response.
4. If the client is banned then return a response that it is banned from the service, else continue to step 5.
5. Decide if the local TPSI is able to serve the client's request or decide to pass the request to a neighbor (based on the load and topology information received from the other TPSIs ; basically, another TPSI from the whole system with load below threshold is chosen randomly among all the TPSIs with loads below threshold).
6. If the request is executed on the local machine then execute it, otherwise wait for it to be executed on another machine.
7. Send the response back to client.

3.2.2. TPSI - TPSI job request

This type of request occurs when the TPSI that received a client job is full and tries to find a "free" TPSI to submit the job to it (a TPSI is "free" if its CPU load is below a predefined threshold). The submitter TPSI acts like a client for the TPSI that receives the request. The only difference is that the TPSI that receives the request doesn't verify the banned status with the Central Supervisory Unit (the submitted TPSI is trusted). The processing stages of this type of request are:

1. Receive the Request (get the request and the Python script).
2. Decode the Request (decide the request type and what logical flow to follow).
3. Decide if the local TPSI is able to serve the client's request or decide to pass the request to a neighbor.
4. If the job is executed on the local machine then execute it, otherwise wait for it to be executed on another machine.
5. Send the response back to client.

3.2.3. TPSI - TPSI IP Multicast communication

Due to their geographical proximity, all the TPSIs are part of the same IP Multicast group. Periodically, each TPSI broadcasts information regarding its load and regarding its neighbors (topology). This information is used by the TPSIs in order to decide where to send a client job when the current TPSI is too loaded.

The Central Supervisory Unit is part of the same IP Multicast group, so it is involved in all

