A Cloud-Based architecture for the Cherenkov Telescope Array observation simulations.

Optimisation, design, and results

cherenkov telescope array

M. Landoni[1], P. Romano[1], S. Vercellone[1], J. Knodlseder[2], A. Bianco[1], F. Tavecchio[1] and A. Corina[3]

[1]*INAF, Osservatorio Astronomico di Brera, Via E. Bianchi 46 I-23807 Merate (LC) – ITALY*
[2]*Institut de Recherche en Astrophysique et Planetologie, 9 avenue Colonel-Roche, 31028, Toulouse Cedex 4, France*
[3]*Amazon Web Services Inc. - P.O. Box 81226, Seattle, WA 98108-1226*

ABSTRACT

Simulating and analyzing detailed observations of astrophysical sources for very high energy (VHE) experiments, like the Cherenkov Telescope Array (CTA), can be a demanding task especially in terms of CPU consumption and required storage. We propose an innovative cloud computing architecture based on Amazon Web Services (AWS) aiming to decrease the amount of time required to simulate and analyze a given field by distributing the workload.

Supported by:

The authors acknowledge contribution from the grant INAF CTA–SKA, 'Probing particle acceleration and γ-ray propagation with CTA and its precursors' (PI F. Tavecchio).

INTRODUCTION

CTA is currently in the scientific assessment phase of simulating feasibility and scientific return of potential astrophysical targets which, in turn, can be used to determine future observing plans that maximize the overall payoff along the whole CTA lifetime. This often implies episodic highly CPU-intensive simulations that are performed on specific science projects within broader topics on very short timescales. Under these conditions, it is *generally not cost effective to purchase set up, and maintain a large enough cluster of computers* to perform the task, and it may be cheaper to buy CPU time (and all correlated services of moving and storing large amounts of data) in a cloud platform.

WHAT WE USE

We combine the Docker Platform to distribute the software and Elastic Cloud Computing (EC2) from Amazon Web Services to distribute the workload and achieve simulation results quickly.

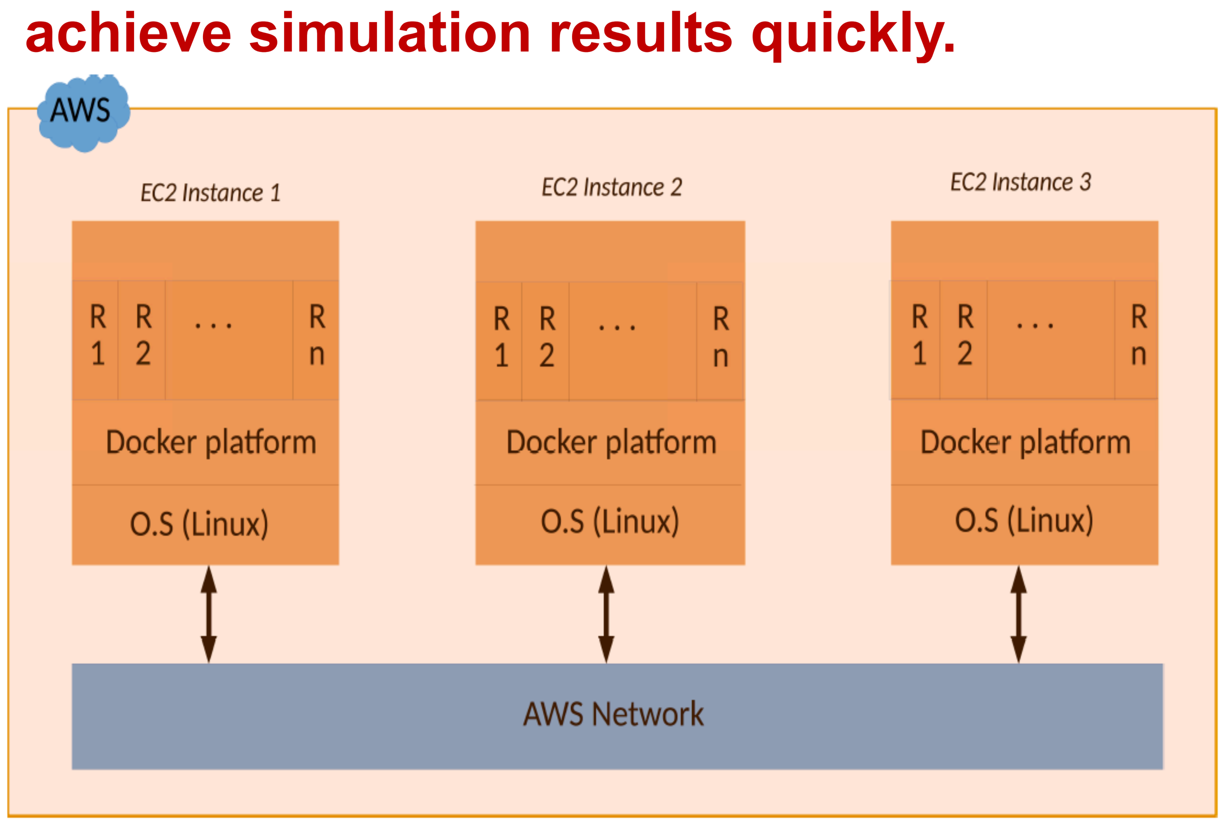

Fig. 1: The *ctools* suite (Knodlseder et al., 2016) is containerized through a Docker in each EC2 instance allowing to run a large number of realizations simultaneously.

ACKNOWLEDGEMENTS

We gratefully acknowledge financial support from the agencies and organizations listed here:
www.cta-observatory.org/consortium_acknowledgments

THE FULL ARCHITECTURE

The distributed computation starts by uploading in the storage S3 (step 1) a *tarball* file containing the data of the source to simulate and the realization list. Then a trigger is raised (step 2) and an Amazon Lambda function pushes the realization list into the distributed first-in first-out queue SQS (step 3a) and creates a set of EC2 instances (step 3b). This allows the distribution of realisations among many computers. When each EC2 instance is online, it automatically pushes messages (that contain information necessary to run a realization) from the SQS queue and starts Docker containers where *ctools* run.

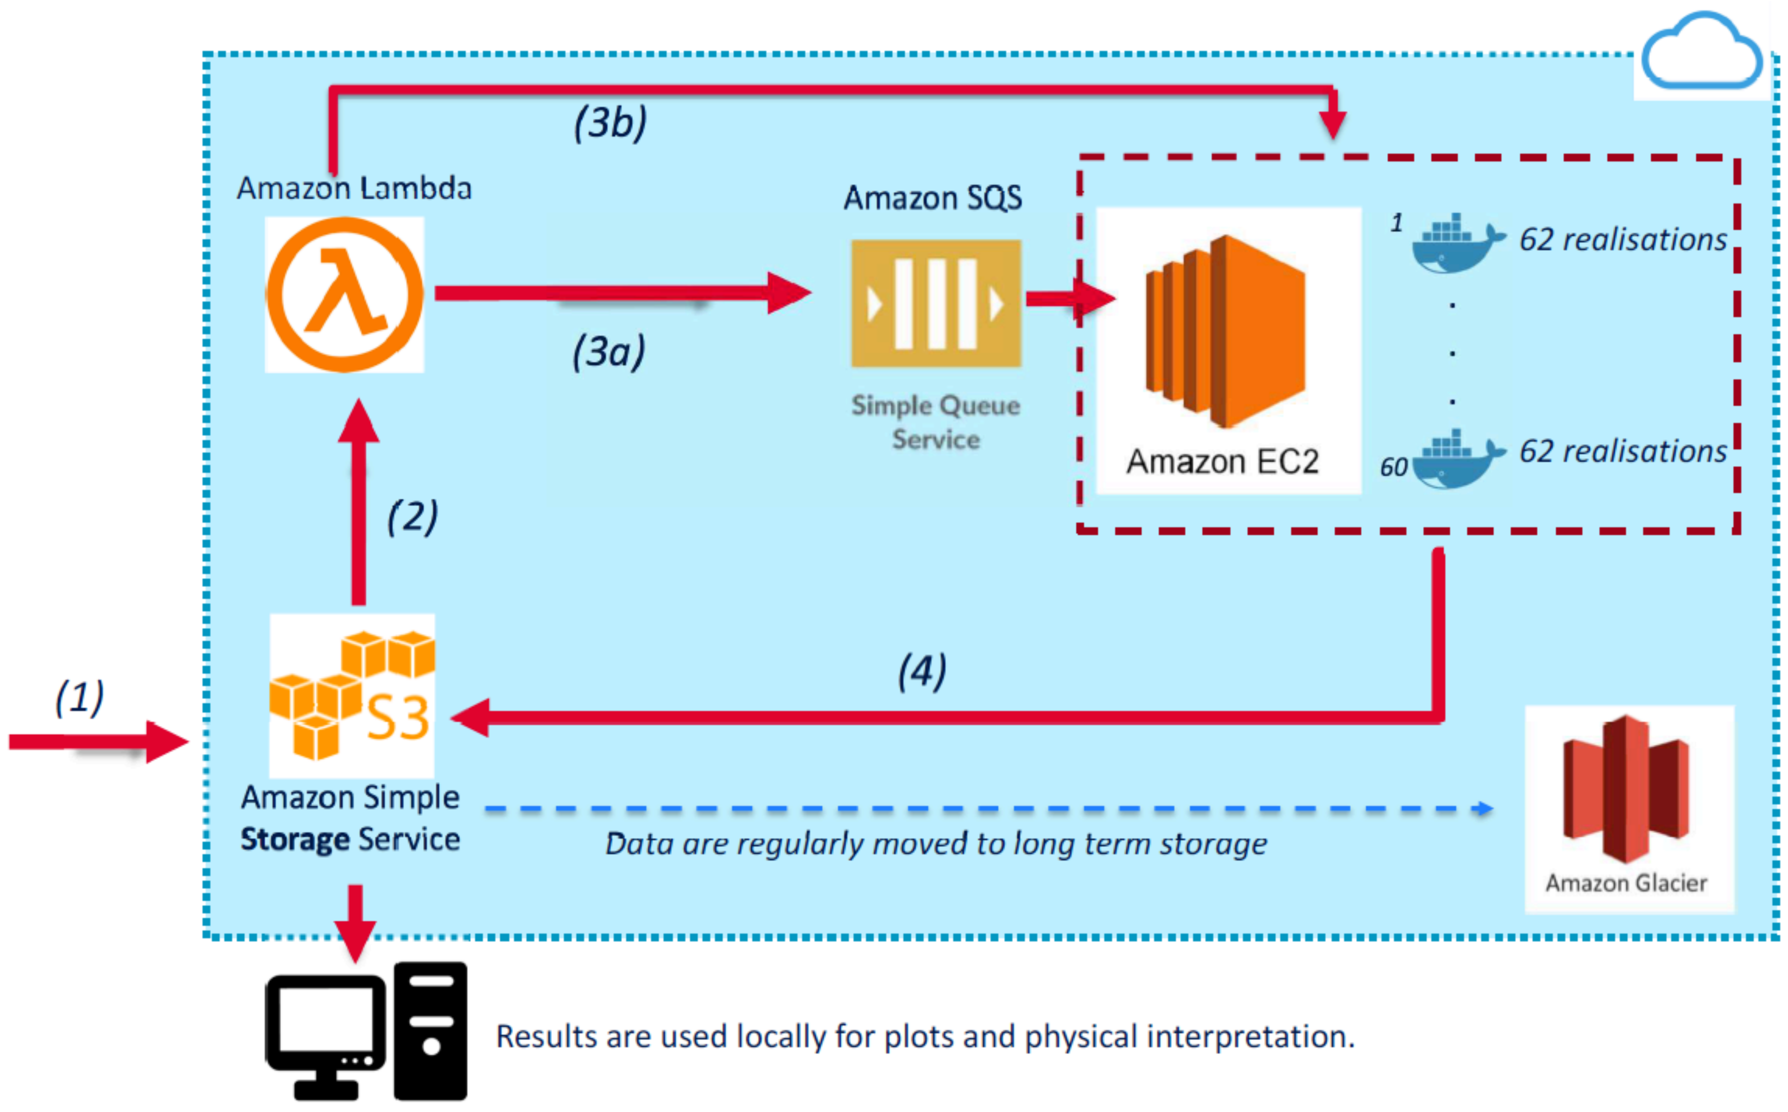

Fig. 2: The cloud-based AWS architecture for parallel and distributed CTA simulations.

Then, the EC2 instance waits for all containers to finish their computation and the outputs, temporarily saved on the local instance storage, are saved onto S3 (step 4). Processed data can then be recovered from S3 for subsequent analysis and physical interpretation of the results.

RESULTS

We find that, by using AWS, we can run our simulations more than 2 orders of magnitude faster than by using a general purpose workstation for the same cost. We suggest to consider this method when observations need to be simulated, and analyzed within short timescales.

Costs (USD)	Local	AWS
Simulations	41.5	1.3
Storage	2.0[a]	5.0[b]
Maintenance	2.5	–
Total	46	6.3
Run Times (hr)	172	0.5
Scaling Factors	Local	AWS
Galactic Background	×5–10	×5–10
Energy dispersion	×5–10	×5–10

References

Landoni M. et al., 2018, ApJS, submitted
Knodlseder J. et al., 2016, A&A, 593, A1