Serverless Services for Scientific Cloud Computing

IberGrid 2022 Conference

Germán Moltó - gmolto@dsic.upv.es
Sebastián Risco - srisco@i3m.upv.es
Vicent Giménez - vigial@posgrado.upv.es
Miguel Caballer - micafer@i3m.upv.es

Agenda

- Motivation for Serverless Computing
- SCAR
- OSCAR
- MARLA
- TaSCaaS
- Summary

UNIVERSITAT POLITÈCNICA DE VALÈNCIA

Motivation: Head Count

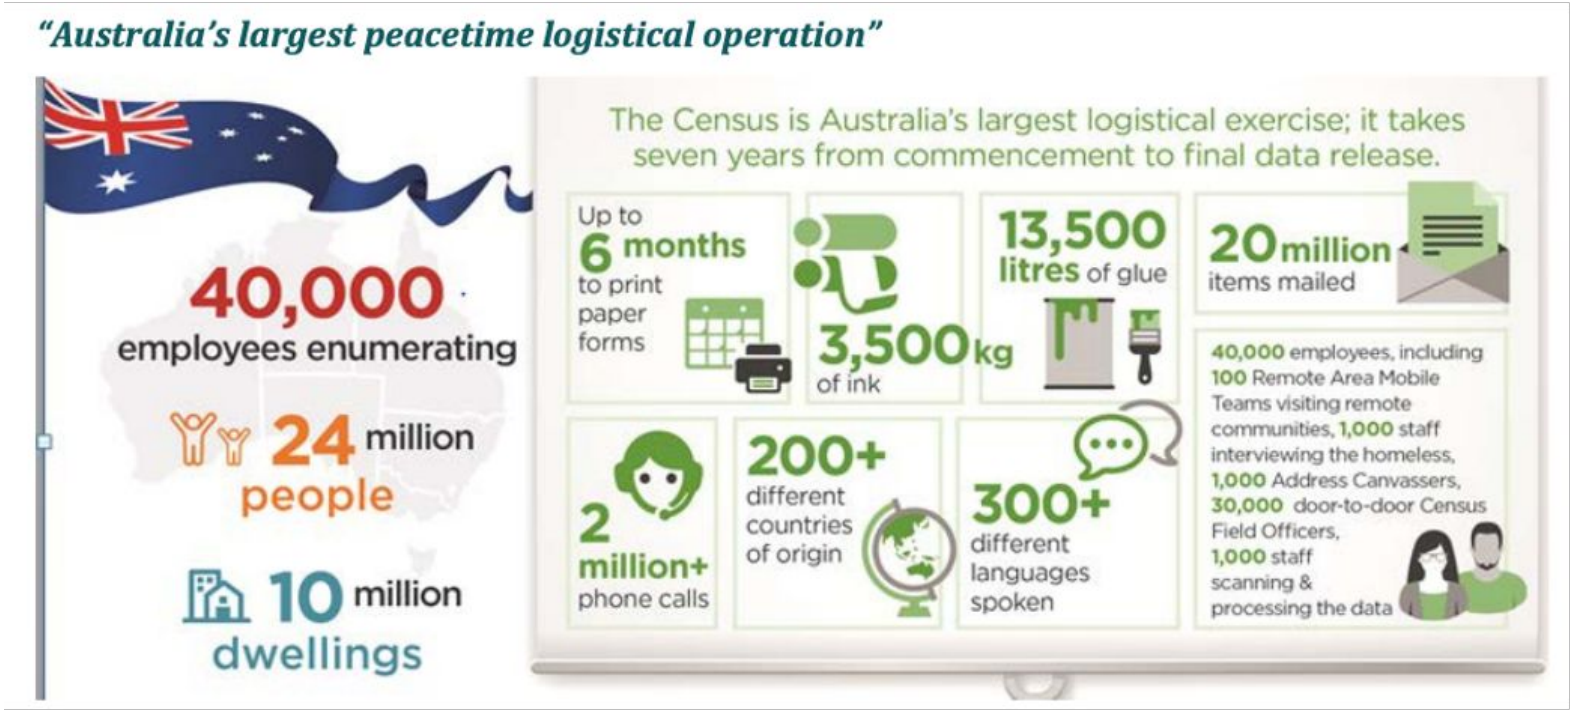

Motivation: Trusting Your Partners

- The Australian Bureau of Statistics, through open tender, awarded IBM a $9.6M a contract to implement an eCensus solution for 2016.
- ABS wisely tendered for services to “Perform Load Testing” ($469K out of which $325K was spent on software licenses).

CN ID: CN2641301
Agency: Australian Bureau of Statistics
Publish Date: 27-Oct-2014
Category: Software maintenance and support
Contract Period: 1-Oct-2014 to 31-Oct-2016
Contract Value (AUD): $9,606,725.00
Description: Design, development and implementation of eCensus Solution 2016

Procurement Method: Limited tender
Confidentiality - Contract: No
Confidentiality - Outputs: No
Consultancy: No
Agency Reference ID: ABS2014.105

Supplier Details

Name: IBM Australia Ltd
Postal Address: 8 Brisbane Ave
Town/City: Barton
Postcode: 2600
State/Territory: ACT
Country: AUSTRALIA
ABN: 79 000 024 733

Motivation: A Story in Three Acts

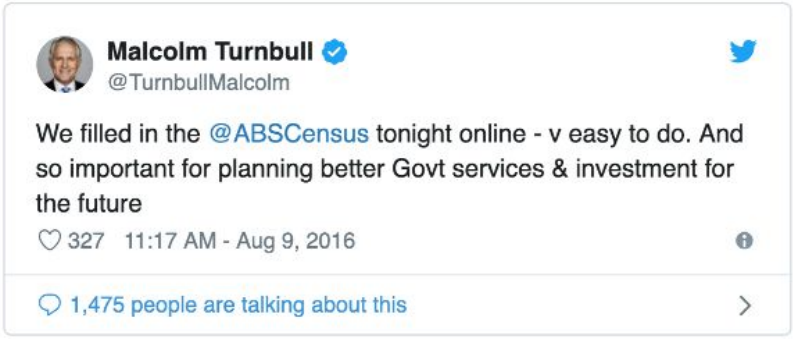

https://twitter.com/TurnbullMalcolm/status/762940763801989121

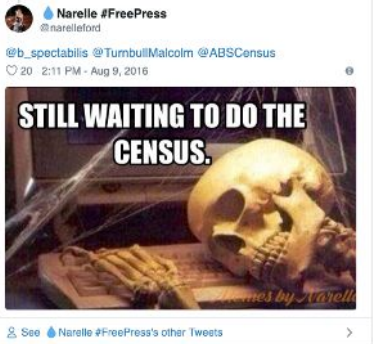

https://twitter.com/narelleford/status/762984702915465216

https://twitter.com/ABSStats/status/762961251764805633

Given that millions of Australians can play Pokemon Go at once and it doesn't crash is a good reason to outsource the census to Nintendo
— Tim Beshara (@Tim_Beshara) August 9, 2016

Motivation: Official vs Unofficial

- Official Statement (13/10/2016) from the Office of Cyber Security Special Adviser:
  - [...] although the site withstood an initial DDoS attack and was coping with over 7,000 census forms a minute, a second and third attack took it down
- Critics: The system was believed to have been built on IBM WebSphere and run on IBM Softlayer (on-premises Cloud) instead of on a public Cloud.

Motivation: A Surprising Turn of Events

- A couple of students, without prior experience in AWS, developed a serverless system over a weekend supporting 4 times the workload used to test IBM's system for ~~$500~~ $30

https://eftm.com/2016/08/how-two-uni-students-built-a-better-census-site-in-just-54-hours-for-500-3075

Motivation: Standing on the Shoulders of Giants

- How could these be possible?
- Students had used AWS Lambda, a massively scalable serverless platform for event-driven computing.
- Serverless: Event-driven computation on a computing platform entirely managed by the Cloud provider

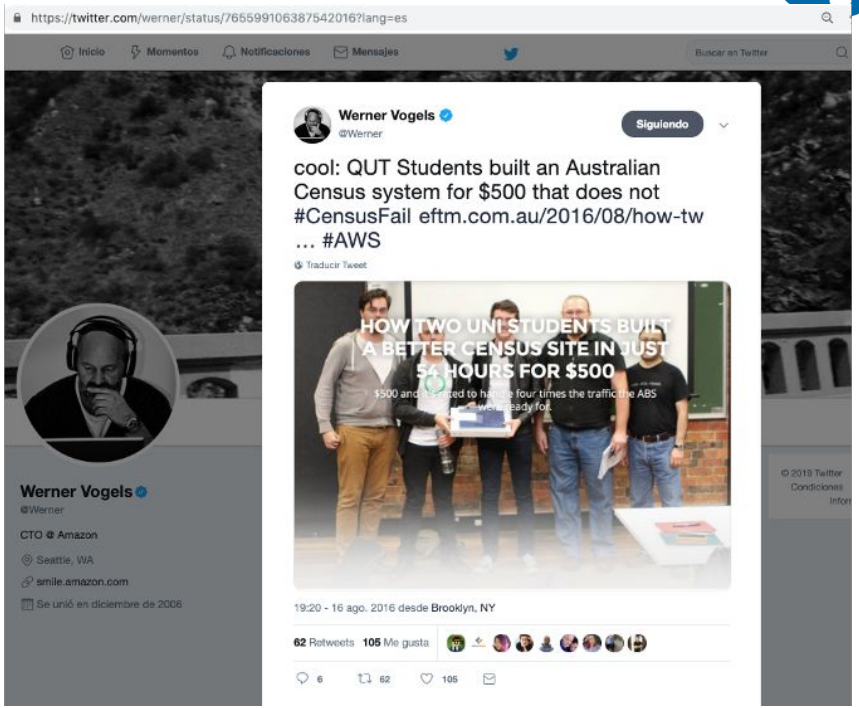

https://twitter.com/werner/status/765599106387542016

Motivation: Long Story Short

- IBM reportedly payed $30M to the Australian government as reports are released from two inquiries into DDoS attacks on census website.
- PwC Australia operated Australian 2021 Digital Census on (quick poll):

Serverless Computing

- Event-driven computing on highly-elastic services with fine-grained billing managed by the Cloud provider.

Storage
-
Amazon
S3

Compute
-
AWS
Lambda

Database
-
Amazon
DynamoDB

Gateways
-
Amazon API
Gateway

Queues
-
Amazon
SQS

Messaging
-
Amazon
SNS

Internet of Things
-
AWS IoT

Streaming Analytics
-
Amazon Kinesis

User Management
-
Amazon
Cognito

Machine Learning
-
Amazon Machine
Learning

SCAR - https://github.com/grycap/scar

- Framework to execution Docker-based applications in AWS Lambda
  - Highly-parallel event-driven file-processing serverless applications that execute on customized runtime environments provided by Docker containers run on AWS Lambda (thanks to uDocker)
  - Pioneered the usage of Docker containers in AWS Lambda since 2017 (native support in AWS Lambda introduced end of 2020, now available in SCAR)
  - Featured in the CNCF Cloud Native Interactive Landscape (>500 stars in GitHub): https://landscape.cncf.io/serverless?selected=scar
- Integrated with API Gateway for HTTP-based scalable endpoints
- Integrated with AWS Batch for cloud-bursting into scalable virtual HPC-based clusters (even with GPU support).

UDOCKER
BE ANYWHERE

SCAR - Architecture

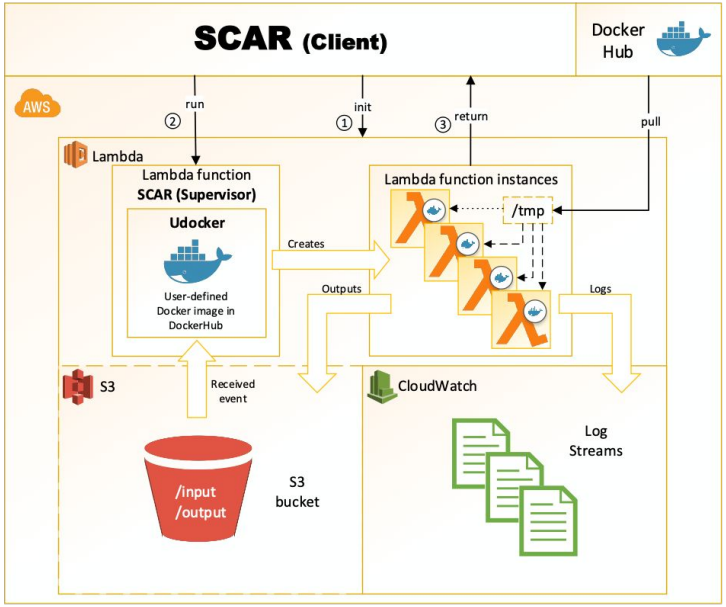

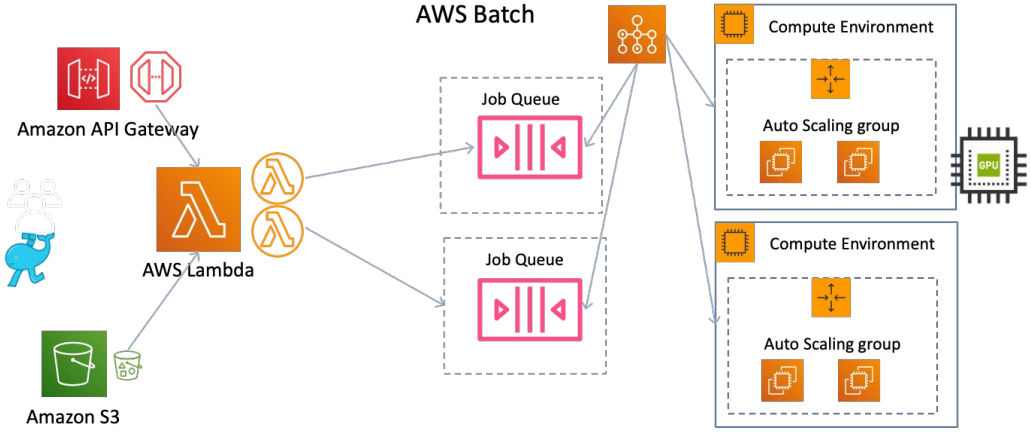

Pérez, A., Moltó, G., Caballer, M., & Calatrava, A. (2018). Serverless computing for container-based architectures. *Future Generation Computer Systems*, *83*, 50–59. https://doi.org/10.1016/j.future.2018.01.022

SCAR - Sample Workflow for Multimedia Processing

- Lambda functions for audio processing
- AWS Batch jobs for video processing

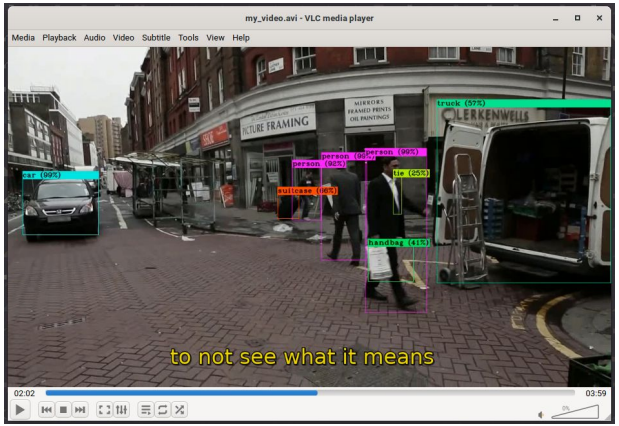

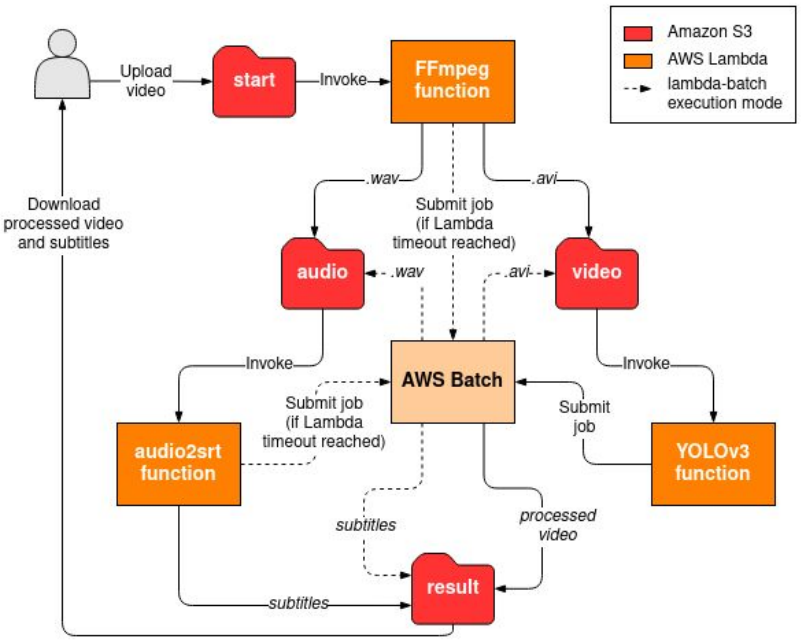

Risco, S., & Moltó, G. (2021). GPU-Enabled Serverless Workflows for Efficient Multimedia Processing. *Applied Sciences*, *11*(4), 1438. https://doi.org/10.3390/app11041438

OSCAR - https://oscar.grycap.net

- Open Source Serverless Computing for Data-Processing Applications (OSCAR)
  - Serverless computing for Docker-based computationally-intensive applications on elastic Kubernetes clusters deployed on multi-Clouds.
- Mimics the event-driven computational paradigm of SCAR but for on-premises (or public) Clouds.

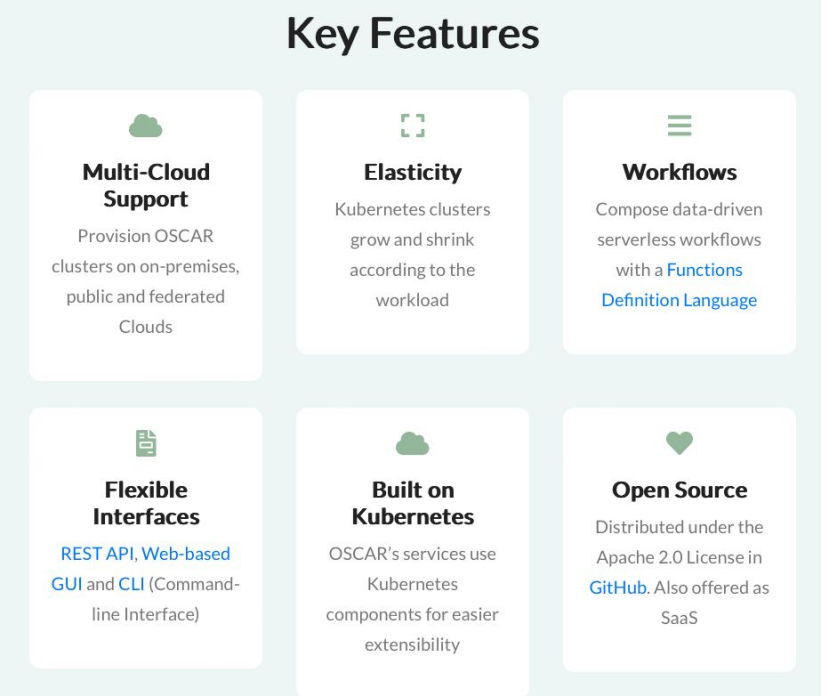

OSCAR - Architecture

- Dynamic provisioning of Kubernetes clusters on multiple Clouds thanks to the Infrastructure Manager (IM) - https://im.egi.eu
- Horizontally scalable Kubernetes clusters thanks to CLUES - https://github.com/grycap/clues

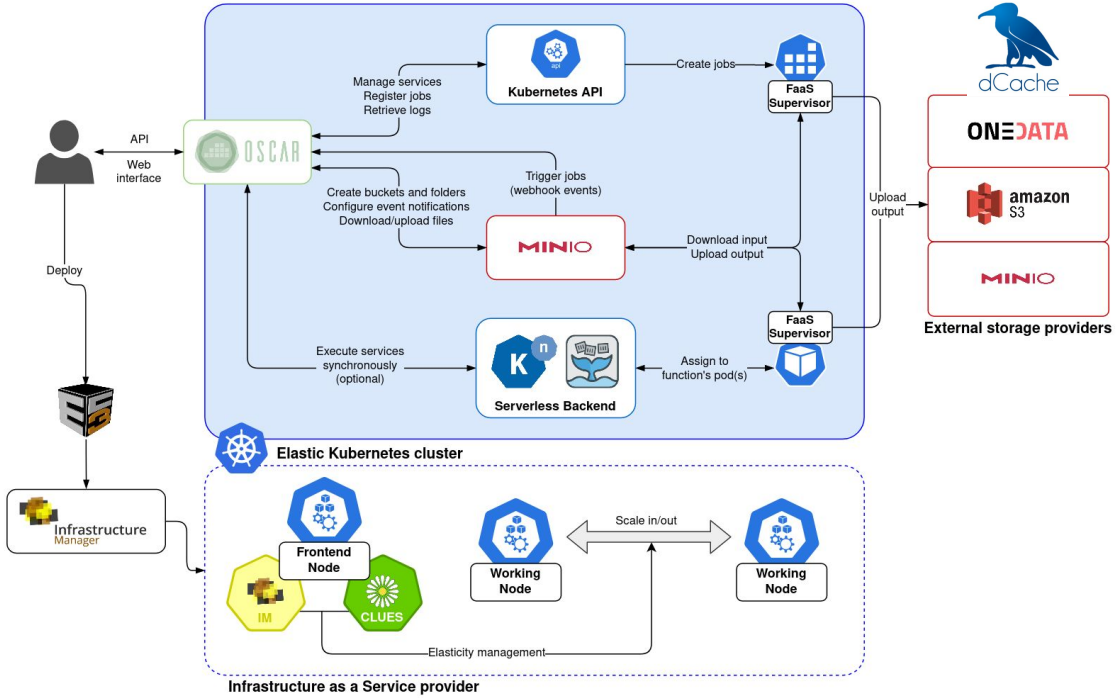

OSCAR - A Sample Use Case

- Mask Detection usage
- Combination of a cluster of Raspberry PIs (edge) and dynamically provisioned resources from EGI for AI inference along the computing continuum

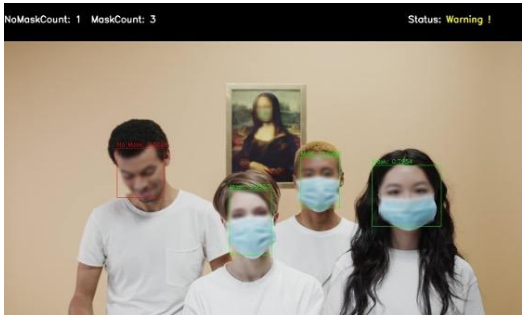

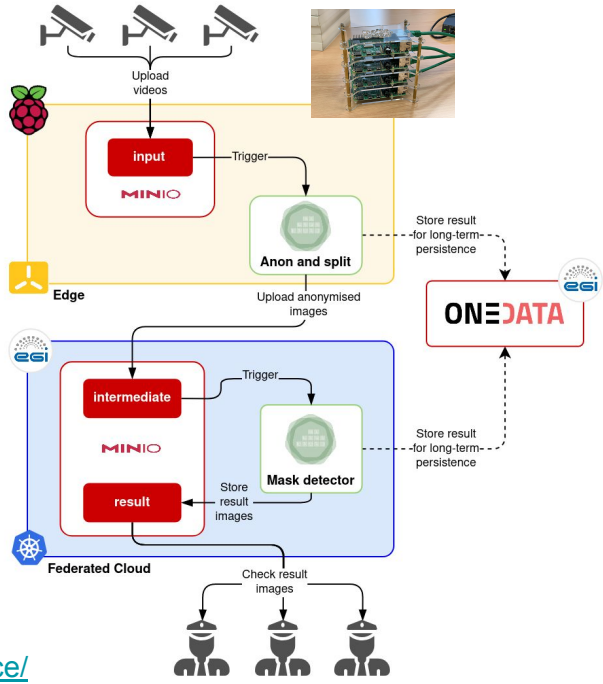

https://oscar.grycap.net/blog/post-oscar-cloud-to-edge-approach-for-edge-ai-inference/

OSCAR + SCAR + EGI

- Functions Definition Language (FDL)

```
functions:
  oscar:
  - oscar-test:
      name: plants
      memory: 2Gi
      cpu: '1.0'
      image: grycap/oscar-theano-plants
      script: plants.sh
      input:
      - storage_provider: minio.default
        path: example-workflow/in
      output:
      - storage_provider: minio.default
        path: example-workflow/med
  - oscar-test:
      name: grayify
      memory: 1Gi
      cpu: '1.0'
      image: grycap/imagemagick
      script: grayify.sh
      input:
      - storage_provider: minio.default
        path: example-workflow/med
      output:
      - storage_provider: minio.default
        path: example-workflow/res
      - storage_provider: onedata.my_onedata
        path: result-example-workflow

storage_providers:
  onedata:
    my_onedata:
      oneprovider_host: my_provider.com
      token: my_very_secret_token
      space: my_onedata_space
```

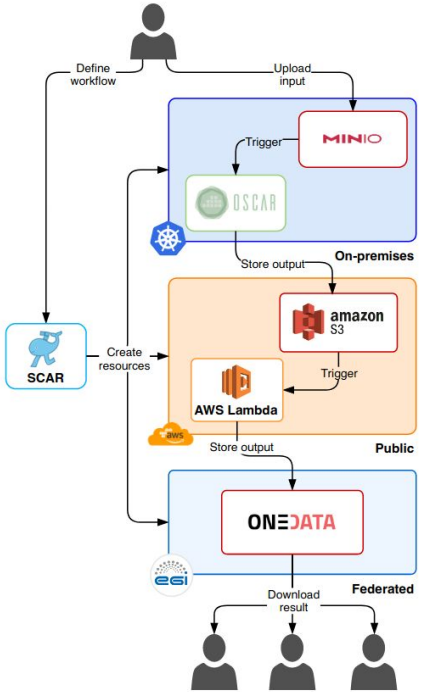

OSCAR - Interfaces

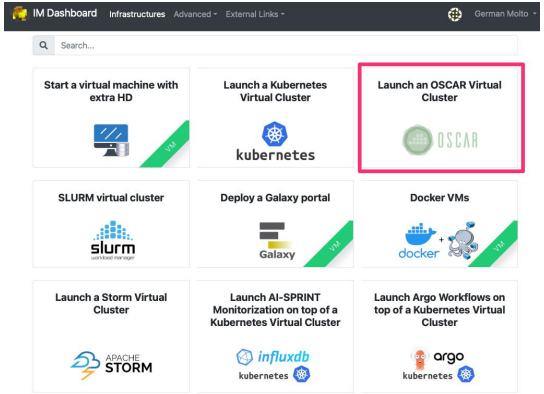

Deployment:
Infrastructure Manager (IM) Dashboard:
https://im.egi.eu

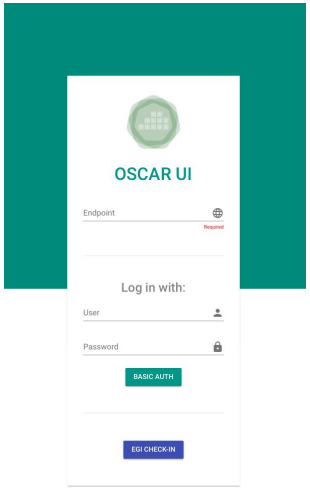

Web-based UI:
https://ui.oscar.grycap.net

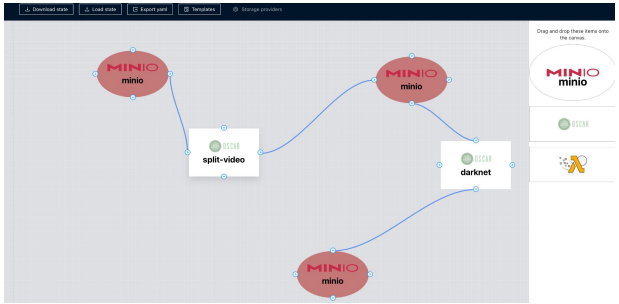

Composer:
https://composer.oscar.grycap.net

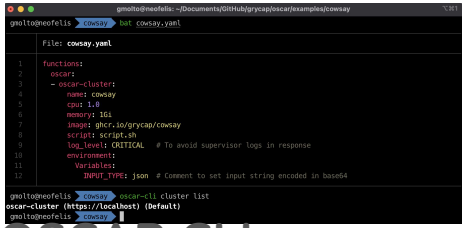

OSCAR CLI:
github.com/grycap/oscar-cli

MARLA - https://github.com/grycap/marla

- Deploy a serverless MapReduce processing engine on AWS Lambda.
- Files uploaded to Amazon S3 trigger the execution of the (parallel invocation of the) functions to concurrently process the dataset.

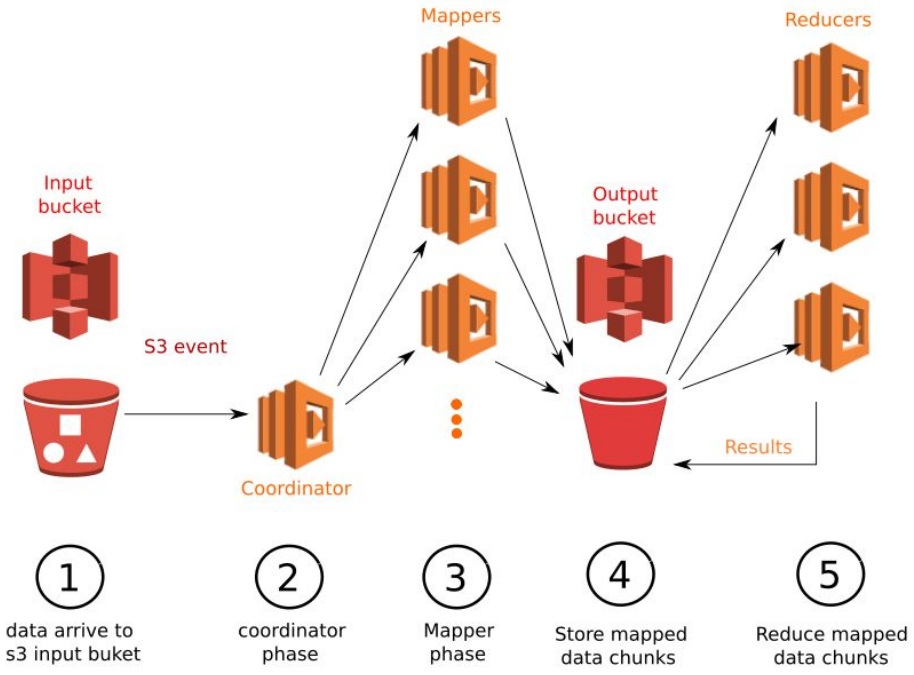

V. Giménez-Alventosa, G. Moltó, and M. Caballer, “A framework and a performance assessment for serverless MapReduce on AWS Lambda,” *Futur. Gener. Comput. Syst.*, vol. 97, pp. 259–274, Aug. 2019, doi: 10.1016/j.future.2019.02.057

MARLA

- The experimental results unveiled that serverless platforms provide inhomogenous computing power that impacts coupled-computing executions of parallel jobs.

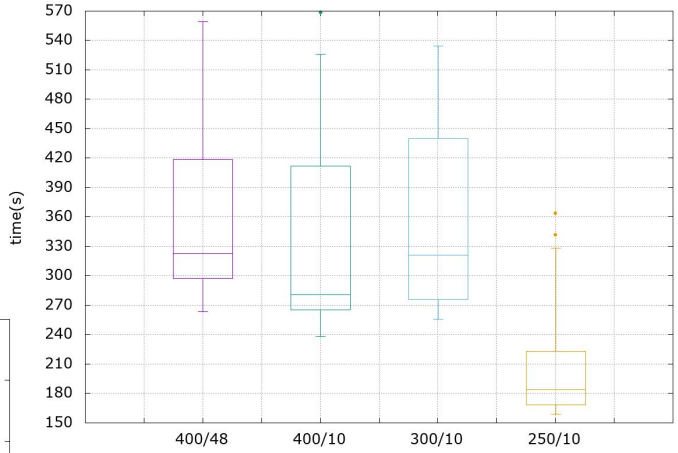

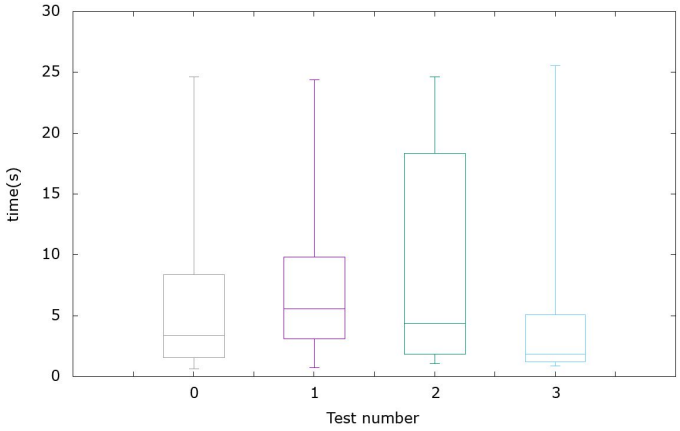

V. Giménez-Alventosa, G. Moltó, and M. Caballer, “A framework and a performance assessment for serverless MapReduce on AWS Lambda,” *Futur. Gener. Comput. Syst.*, vol. 97, pp. 259–274, Aug. 2019, doi: 10.1016/j.future.2019.02.057

TaScaaS - https://github.com/grycap/tascaas

- Task Scheduler As A Service (TaScaaS) provides a complete serverless service to schedule and distribute High Throughput Computing (HTC) jobs among computational infrastructures.

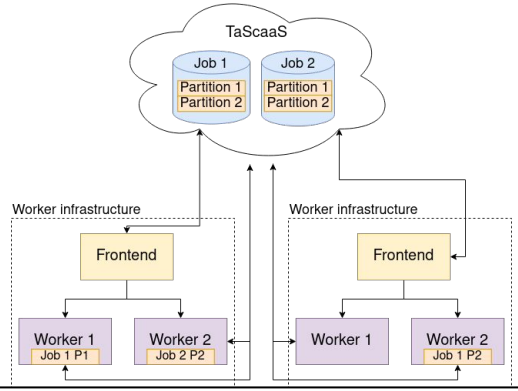

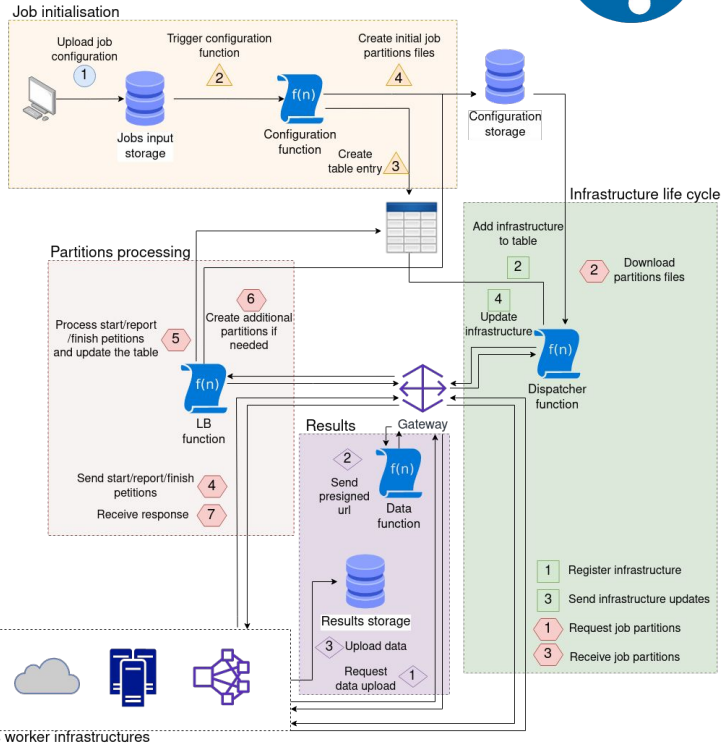

V. Gimenez-Alventosa, G. Molto, and J. D. Segrelles, “TaScaaS: A Multi-Tenant Serverless Task Scheduler and Load Balancer as a Service,” IEEE Access, vol. 9, pp. 125215–125228, 2021, doi: 10.1109/ACCESS.2021.3109972.

Conclusions

- Event-driven computing allows to perform computing in response to events (such as file uploads) on a serverless platform which provides automated elasticity for dynamic resource provisioning.
- SCAR executes generic applications on AWS Lambda and automated extension to AWS Batch to allow elastic CPU/GPU batch computing on the Cloud.
- OSCAR implements the event-driven computing model of SCAR in on-premises Clouds, integrated with EGI services (EGI DataHub and EGI Federated Cloud).
- Innovative services can be developed on the foundations of serverless computing platforms, such as MARLA and TaSCaaS.

Contact

@gmolto

Germán Moltó

gmolto@dsic.upv.es

Instituto de Instrumentación para la Imagen Molecular (I3M)

Universitat Politècnica de València

Grant PID2020-113126RB-I00 funded by MCIN/AEI/10.13039/501100011033.

Project PDC2021-120844-I00 funded by MCIN/AEI/10.13039/501100011033 funded by the European Union NextGenerationEU/PRTR

Part of this work was supported by the project AI-SPRINT ‘‘AI in Secure Privacy-Preserving Computing Continuum’’ that has received funding from the European Union’s Horizon 2020 Research and Innovation Programme under