



Grid y Computación
de Altas Prestaciones

GRyCAP

Instituto de Instrumentación para
Imagen Molecular
Universitat Politècnica de València
Spain

SERVERLESS COMPUTING FOR DATA-PROCESSING ACROSS PUBLIC AND FEDERATED CLOUDS

Sebastián Risco, Alfonso Pérez, Miguel Caballer, Germán Moltó

IBERGRID 2019

September 23-26, Santiago de Compostela, Spain



UNIVERSITAT
POLITÈCNICA
DE VALÈNCIA





- Motivation
- Goals
- Components
- Architecture
- Use case
- Conclusions
- Future work



- Public Cloud Serverless services are evolving from the initial FaaS approach to also embrace the execution of containerised applications.
 - AWS Fargate, Google Cloud Run, **AWS Batch**.
- Scientific applications may require specific resources (large amount of memory or CPUs, accelerated devices, etc).
 - Private or Federated Clouds not always fulfil the requirements.
- Federated storage for data persistence remains suitable for scientific applications.

GOALS



- Execute hybrid Serverless workloads using public Clouds for computing and federated storage for data persistence.
 - AWS services to run containerised data-processing applications and EGI DataHub as a storage back-end.
- Automatically delegate longer executions, as well as those requiring specialised hardware (GPUs), to AWS Batch.
- Demonstrate the feasibility of this approach through a use case in video processing.
 - GPU-based computing in the public Cloud to dramatically accelerate object recognition.



- **AWS Lambda:**

- Public Functions as a Service (FaaS) platform.
- No infrastructure provision or configuration management
- Automated elasticity.
- Supports Java, Go, PowerShell, Node.js, C#, Python, and Ruby code.
- Function limits: 3008 MB Memory and 15 minutes execution timeout.



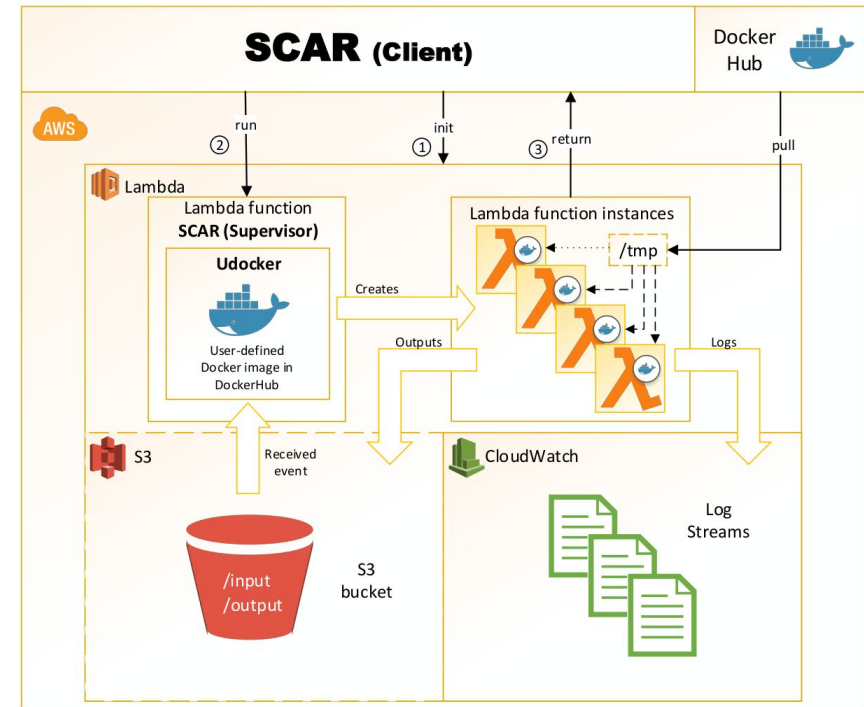
- **AWS Batch:**

- Execute jobs as containerized applications running on Amazon ECS.
- Granular job definitions → specify resource requirements, IAM roles, volumes, GPU access, etc.
- Dynamic compute resource provisioning and scaling.
- No timeout.





- Serverless Container-aware ARchitectures (SCAR):
 - Run containerised applications on AWS Lambda.
 - Defines an event-driven file-processing programming model.
 - **Integrated with AWS Batch in order to support long-running jobs and accelerated computing.**



A. Pérez, G. Moltó, M. Caballer, and A. Calatrava, "Serverless computing for container-based architectures", *Futur. Gener. Comput. Syst.*, vol. 83, pp. 50–59, Jun. 2018.

<https://github.com/grycap/scar>



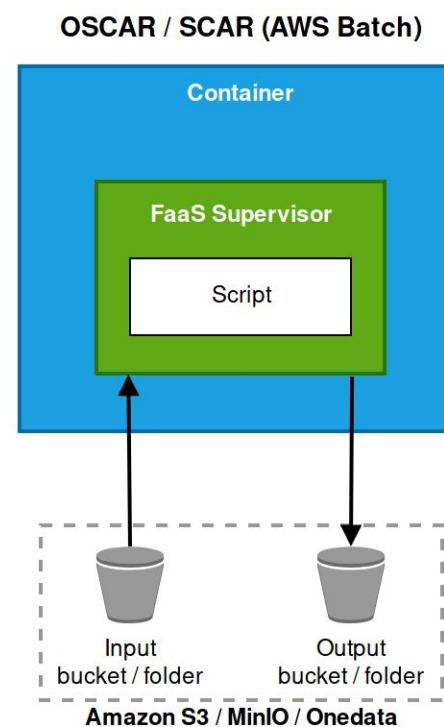
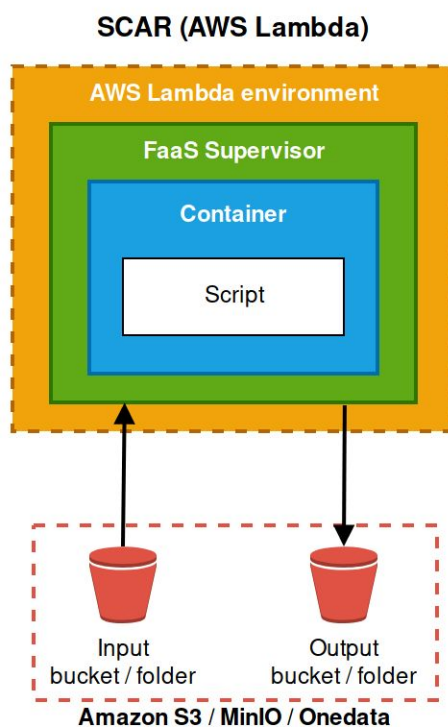
- EGI Data Hub:
 - Service to make data discoverable and available in an easy way across all EGI federated resources, based on Onedata:
 - High-performance data management solution that offers unified data access across globally distributed environments and multiple types of underlying storage.
 - Allows users to share, collaborate and perform computations on the stored data easily.



- OneTrigger:
 - Tool to detect Onedata file events in order to trigger a webhook.
 - It can run as a **Serverless function using AWS Lambda and CloudWatch Events.**



- FaaS Supervisor (Core component of SCAR and OSCAR):
 - Manages input and output.
 - Handles the execution of the user-defined script.
 - Loads Docker containers in AWS Lambda environments.
 - **Integrated with Onedata.**

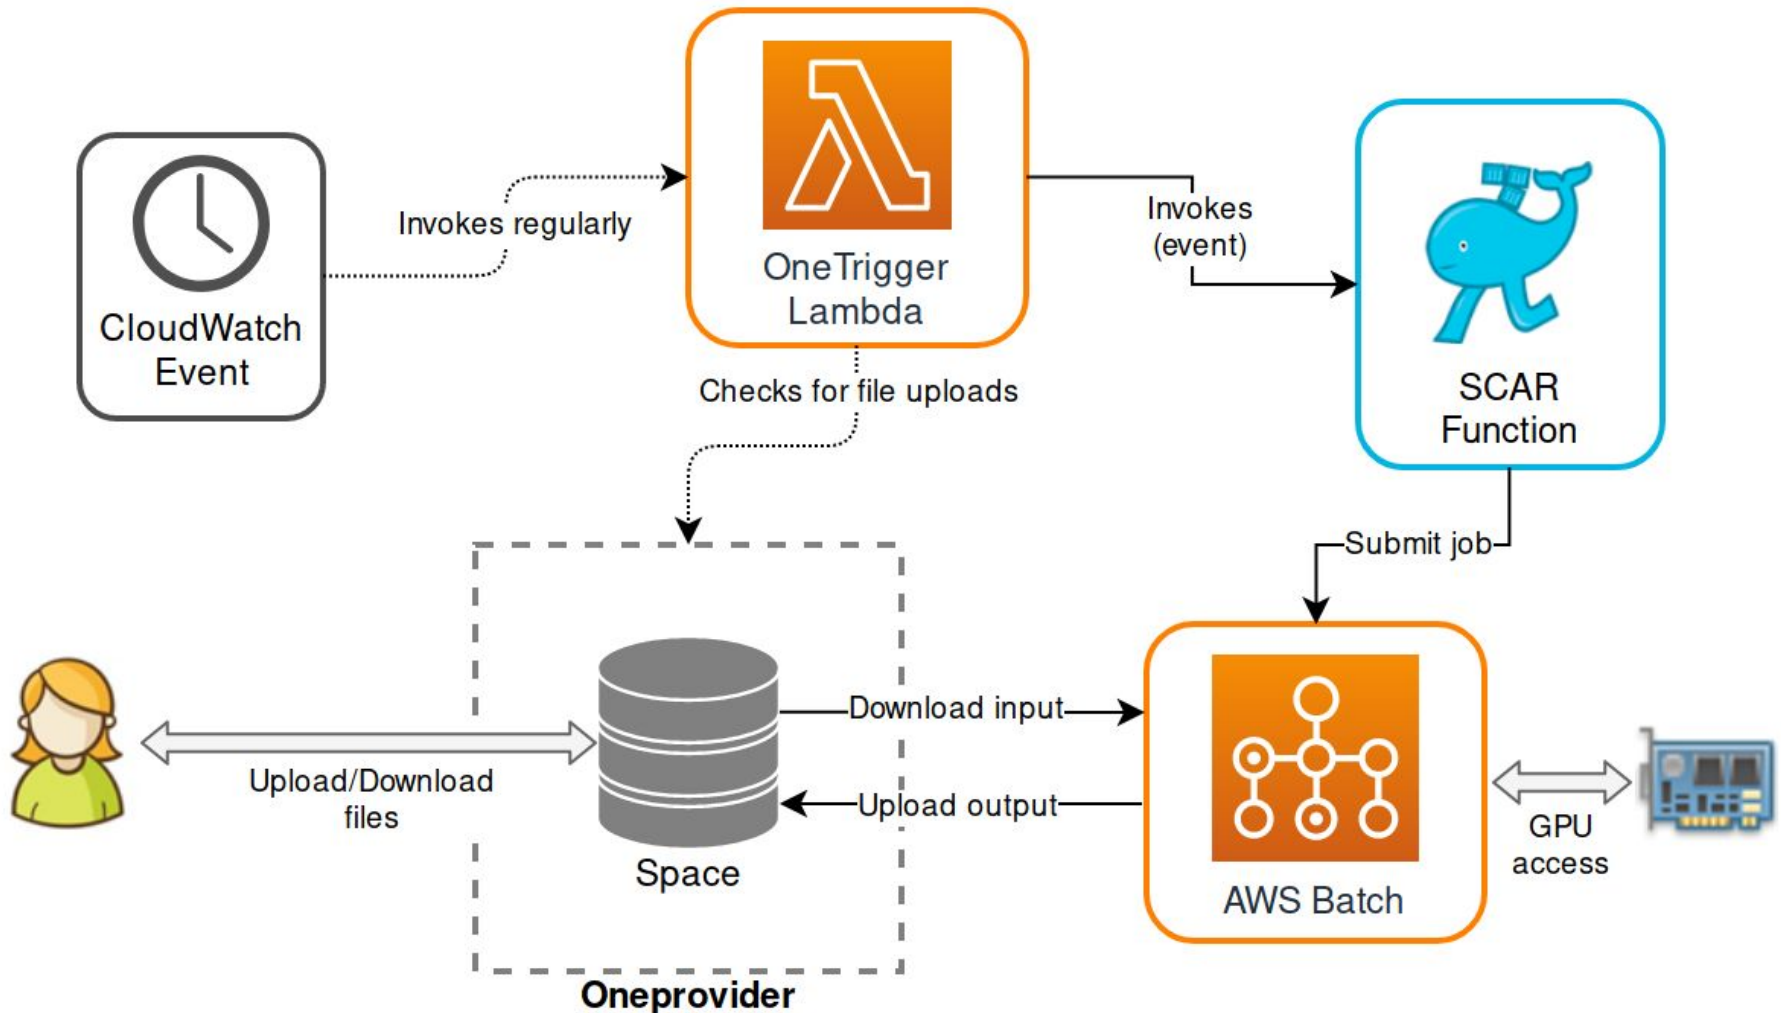


ARCHITECTURE




GRyCAP
Grid y Computación de Altas Prestaciones

www.grycap.upv.es



YOLO (You Only Look Once):

- Real-time object detection system.
- Uses Darknet, an open source neural network framework.
 - Supports CPU and GPU computation.
- Can process images or videos.






Why is GPU recommended for video processing?

- Processing a single image could take few seconds using a CPU.
- If we want the result in images:
 - The video can be split into images.
 - Images can be quickly processed in parallel functions using a Serverless platform (over CPU).
- If we want the result as a video:
 - It has to be processed as a single job.
 - OpenMP can be used to accelerate processing in multi-core CPUs → **It's still very slow.**



CPU vs GPU comparison

Object detection in 1 minute @ 30 FPS video using YOLOv3





- SCAR function definition file

functions:

scar-yolo-video:

image: srisco/yolov3:opencv-cudnn	Docker image
init_script: user-script.sh	User-defined script
s3:	Create input bucket in AWS S3
input_bucket: scar-yolo-video	
api_gateway:	Create HTTP endpoint in AWS API Gateway
name: scar-yolo-video	
execution_mode: batch	Enable AWS Batch mode
batch:	
enable_gpu: true	
compute_resources:	
max_v_cpus: 4	
instance_types:	
- p2.xlarge	AWS Batch configuration
lambda_environment:	
STORAGE_AUTH_ONEDATA_SPACE_1: my-onedata-space	
STORAGE_AUTH_ONEDATA_HOST_1: plg-cyfronet-01.datahub.eui.eu	
STORAGE_AUTH_ONEDATA_TOKEN_1: my-secret-token	
STORAGE_PATH_INPUT_1: scar-yolo-video-in	
STORAGE_PATH_OUTPUT_1: scar-yolo-video-out	Onedata required environment variables

USE CASE



- Integration with EGI DataHub (Onedata)


Lambda > Functions > scar-onetrigger ARN - arn:aws:lambda:us-east-1:974349055189:funcio



scar-onetrigger



Throttle Qualifiers Actions Select a test event Test Save

Configuration Monitoring




▼ Designer



 scar-onetrigger
 Layers (1)

 CloudWatch Events ×
 API Gateway ×

+ Add trigger

 AWS Batch
 Amazon CloudWatch Logs
 Amazon S3

Resources that the function's role has access to appear here

CloudWatch Events

grycap-onetrigger
arn:aws:events:us-east-1:974349055189:rule/grycap-onetrigger
Event bus: **default** Schedule expression: **rate(5 minutes)**

☒ Enabled Delete

USE CASE



ONE DATA
plg-cyfronet-01

sriscs-space > scar-yolo-video-in

FILES	SIZE	MODIFICATION
seq1.avi	16 MiB	2019-08-07 12:17

seq1.avi - VLC media player

Media Playback Audio Video Subtitle Tools View Help

00:01 01:07

USE CASE



aws Services Resource Groups

AWS Batch
Dashboard
Jobs
Job definitions
Job queues
Compute environments

Create job

Job queues

Last loaded: 11:54:05 am 08/07/19

Name	Priority	SUBMITTED	PENDING	RUNNABLE	STARTING	RUNNING	FAILED	SUCCEEDED
scar-plant-classification	1	0	0	0	0	0	0	0
scar-yolo-video	1	0	0	1	0	0	0	0

Compute environments

Last loaded: 11:54:05 am 08/07/19



aws Services Resource Groups

AWS Batch
Dashboard
Jobs
Job definitions
Job queues
Compute environments

Jobs

Jobs are containerized workloads that you submit to AWS Batch.

Search by job ID

Submit job Clone job Cancel job Terminate job

Queue: scar-yolo-video

Status: submitted pending runnable starting running **succeeded** failed < Page 1 >

Job ID	Job name	Array size	Number of nodes	Status	Created at	Started at	Run time
6483bb09-6696-4d8d-a819-08689e55d9a6	scar-yolo-video	--	--	SUCCEEDED	12:18:15 pm 08/07/19	12:21:23 pm 08/07/19	3 minutes

USE CASE



CloudWatch > Log Groups > /aws/batch/job > scar-yolo-video/default/6c67d1d1-83ca-48ff-82d4-dbf4b5f79038



Try CloudWatch Logs Insights

CloudWatch Logs Insights allows you to search and analyze your logs using a new, purpose-built query language. Click [here](#) to experience it. If you want to learn more, read [the AWS blog](#) or visit [our documentation](#).

Expand all ☒ Row ☐

Filter events

all 30s 5m 1h

Time (UTC +00:00)	Message
2019-08-07	
▶ 10:23:55	cvWriteFrame
▶ 10:23:55	Objects:
▶ 10:23:55	FPS:14.2
▶ 10:23:55	cvWriteFrame
▶ 10:23:55	Objects:
▶ 10:23:55	FPS:14.2
▶ 10:23:55	cvWriteFrame
▶ 10:23:55	Objects:
▶ 10:23:55	FPS:14.2
▶ 10:23:55	cvWriteFrame
▶ 10:23:55	Objects:
▶ 10:23:55	FPS:14.1
▶ 10:23:55	Stream closed.
▶ 10:23:55	cvWriteFrame
▶ 10:23:55	input video stream closed.
▶ 10:23:55	closing... closed!output_video_writer closed.
▶ 10:23:55	2019-08-07 10:23:55,657 - supervisor - INFO - Reading output path variables
▶ 10:23:55	2019-08-07 10:23:55,658 - supervisor - INFO - Found 'S3' output provider
▶ 10:23:55	2019-08-07 10:23:55,658 - supervisor - INFO - Searching for files to upload in folder '/tmp/tmpptko8fra'
▶ 10:23:55	2019-08-07 10:23:55,658 - supervisor - INFO - Found the following files to upload: ['/tmp/tmpptko8fra/seq1.avi']
▶ 10:23:55	2019-08-07 10:23:55,658 - supervisor - INFO - Uploading file 'scar-yolo-video/output/f6d58e30-d5d2-45e3-8e38-1549255c0361/seq1.avi' to bucket 'scar-yolo-video'
▶ 10:23:58	2019-08-07 10:23:58,736 - supervisor - INFO - Changing ACLs for public-read for object in bucket 'scar-yolo-video' with key 'scar-yolo-video/output/f6d58e30-d5d2-45e3-8e38-1549255c0361/seq1.avi'
▶ 10:23:58	2019-08-07 10:23:58,821 - supervisor - INFO - Found 'ONEDATA' output provider
▶ 10:23:58	2019-08-07 10:23:58,821 - supervisor - INFO - Searching for files to upload in folder '/tmp/tmpptko8fra'
▶ 10:23:58	2019-08-07 10:23:58,821 - supervisor - INFO - Found the following files to upload: ['/tmp/tmpptko8fra/seq1.avi']
▶ 10:23:58	2019-08-07 10:23:58,822 - supervisor - INFO - Uploading file 'seq1.avi' to 'srisco-space/scar-yolo-video-out'
▶ 10:24:01	2019-08-07 10:24:01,351 - supervisor - INFO - Creating response

No newer events found at the moment. [Retry](#).

USE CASE



ONE DATA
plg-cyfronet-01

srisco-space > scar-yolo-video-out

FILES	SIZE	MODIFICATION
seq1.avi	15.1 MIB	2019-08-07 12:24

public
scar-imagemagick-in
scar-imagemagick-out
scar-plants-in
scar-plants-out
scar-yolo-video-in
scar-yolo-video-out
split-video-in

seq1.avi - VLC media player

Media Playback Audio Video Subtitle Tools View Help

person (98%)
cell_phone (35%)
person (100%)

00:01 01:07

CONCLUSIONS



- Delegating computational jobs to public Cloud providers is convenient for certain cases (even though when private or federated resources are available).
- Serverless allows to reduce costs in longer or accelerated executions.
- Hybrid workflows enable fully leveraging of cloud capabilities in order to run scientific applications.



- Support additional storage back-ends.
- OneTrigger improvements:
 - More efficient file upload checking.
 - Integrate OneTrigger-Lambda with the CLI to automate deployment.
 - Send events to functions directly (without API Gateway).
- Integrate more use cases.
- We are accepting contributions at:

<https://github.com/grycap/scar>

<https://github.com/grycap/faas-supervisor>

<https://github.com/grycap/onetrigger>

CONTACT & ACKNOWLEDGEMENTS



GRyCAP
Grid y Computación de Altas Prestaciones

www.grycap.upv.es

Sebastián Risco - serisgal@izm.upv.es

Alfonso Pérez - alpegon3@upv.es

Miguel Caballer - micafer1@upv.es

Germán Moltó - gmolto@dsic.upv.es

Instituto de Instrumentación para Imagen Molecular

Universitat Politècnica de València

Camino de Vera s/n

46022, Valencia

SPAIN



The authors would like to thank the Spanish "Ministerio de Economía, Industria y Competitividad" for the project "BigCLOE" with reference number TIN2016-79951-R.

This work has been partially funded through the EGI Strategic & Innovation Fund.