Serverless: What's in name for scientific computing?

Germán Moltó – gmolto@dsic.upv.es

IBERGRID 2019, 23-26 September, Santiago de Compostela, Spain

Motivation

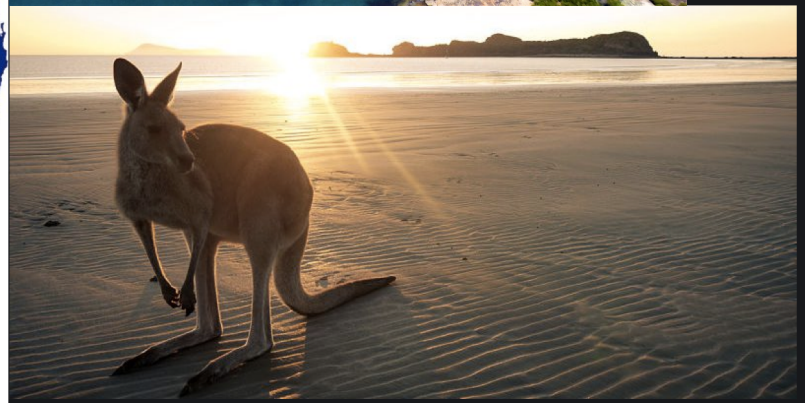

Head count

- Every 5 years, Australians update their census.

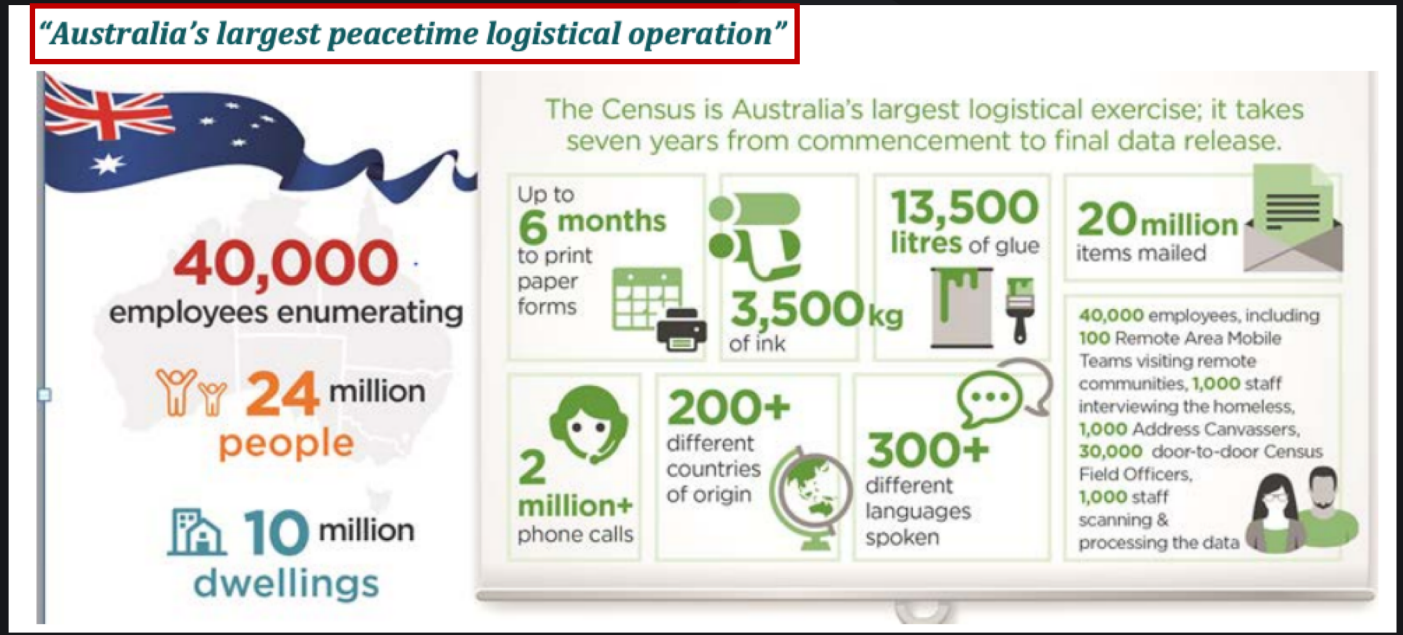

https://apo.org.au/sites/default/files/resource-files/2016/11/apo-nid70705-1232016.pdf

Trusting your Partners

- The ABS*, 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).

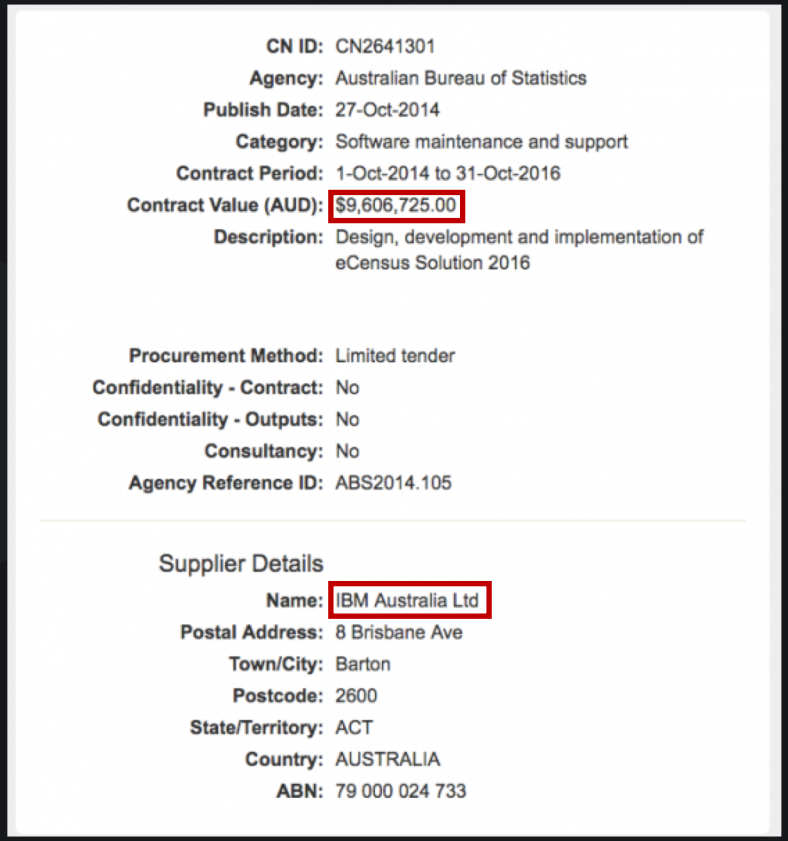

CN ID: CN2641301
Agency: Australian Bureau of Statistics
Publish Date: 27-Oct-2014
Category: Software maintenance and support
Contract Period: 1-Oct-2014 to 31-Oct-2016
Contract Value (AUD): $9,606,725.00
Description: Design, development and implementation of eCensus Solution 2016

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

* Australian Bureau of Statistics

A Story in Three Acts

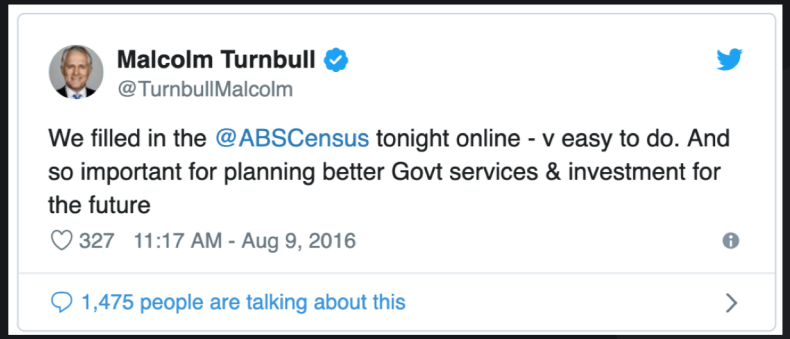

https://twitter.com/TurnbullMalcolm/status/762940763801989121

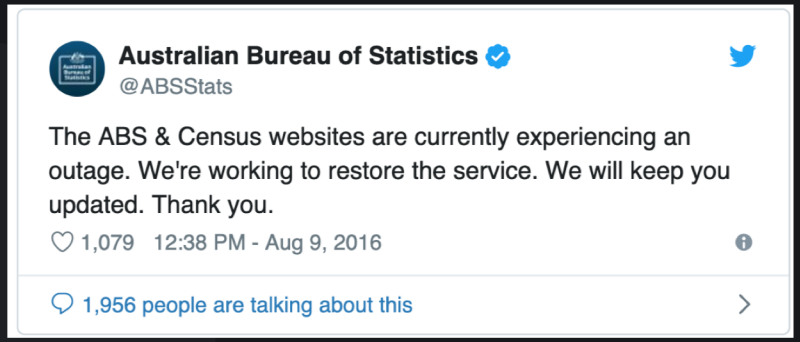

https://twitter.com/ABSStats/status/762961251764805633

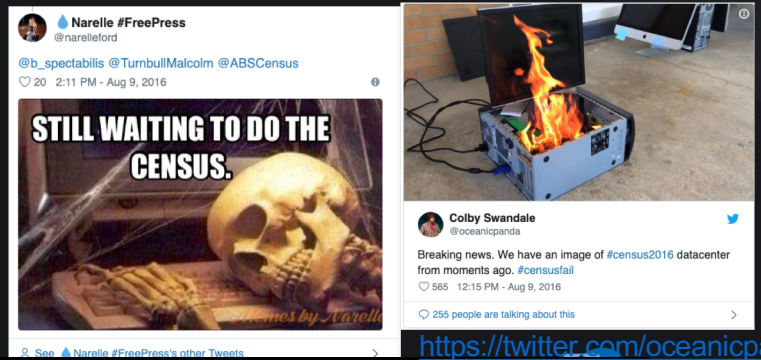

https://twitter.com/narelleford/status/762984702915465216

https://twitter.com/oceanicpanda/status/762955516096094208

> 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

Official vs Unofficial

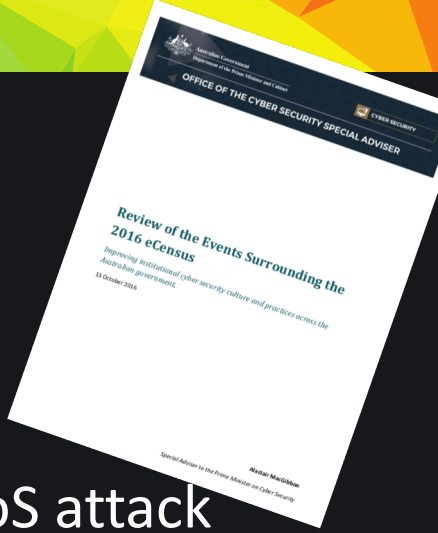

- Official Statement (13/10/2016) from the Office of Cyber Security Special Adviser:
  - […] although the site withstood an initial DDoS attack and was coping with over 7,000 census forms a minute, a second and third attack took it down
- Critics: The system was believed to have been built on IBM WebSphere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

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-500-30752

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.

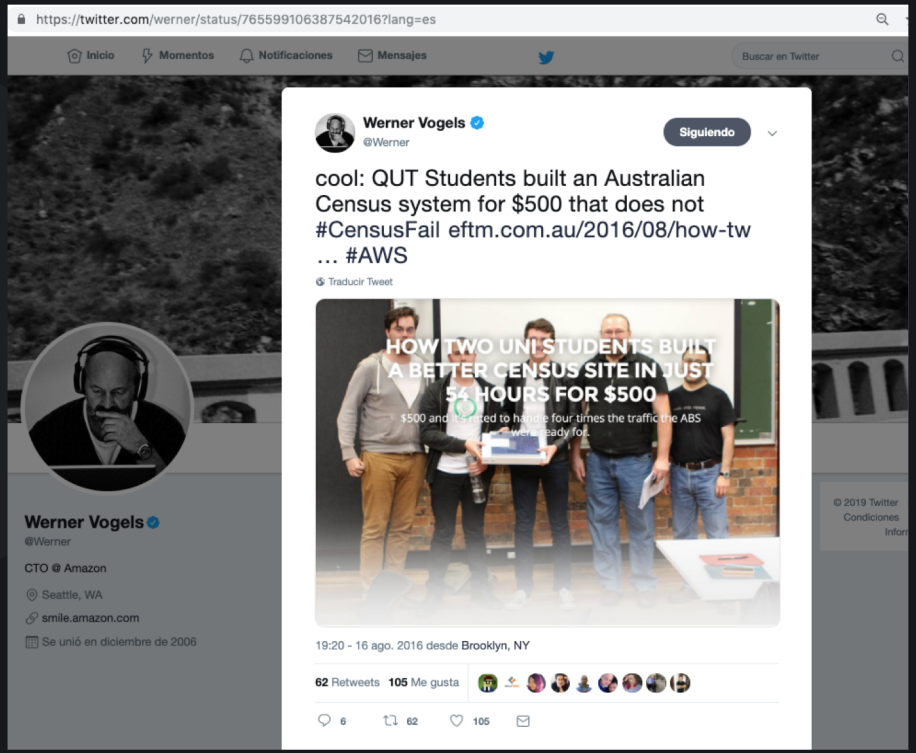

https://twitter.com/werner/status/765599106387542016

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 will operate Australian 2021 Digital Census on (quick poll):

Who's Speaking?

- Germán Moltó - https://www.grycap.upv.es/gmolto
  - Associate Professor at the Universitat Politècnica de València.
  - Researcher in Serverless/Clouds for scientific computing.
  - Participat(ed/ing) in several European Cloud projects:
    - INDIGO-DataCloud, EOSC-HUB, EOSC-Synergy, DEEP Hybrid DataCloud, etc.

Outline of the Talk

1. Motivation and Introduction
2. What is Serverless Computing?
3. Public Serverless service: AWS Lambda
4. Serverless for scientific computing
5. Serverless (on-premises!)
6. Conclusions

User-defined Cloud Services

- User-defined Cloud services require to manage:
  - Data (i.e. *State,* in the shape of files, databases, in-memory values, etc.)
  - Computing (resources and execution environment).
- R*esilient* application : Manage *Replication* and *Distribution* of both data and computing.

Pre-Serverless

- Deploying highly-available applications is far from being a trivial task in the pre-serverless era.

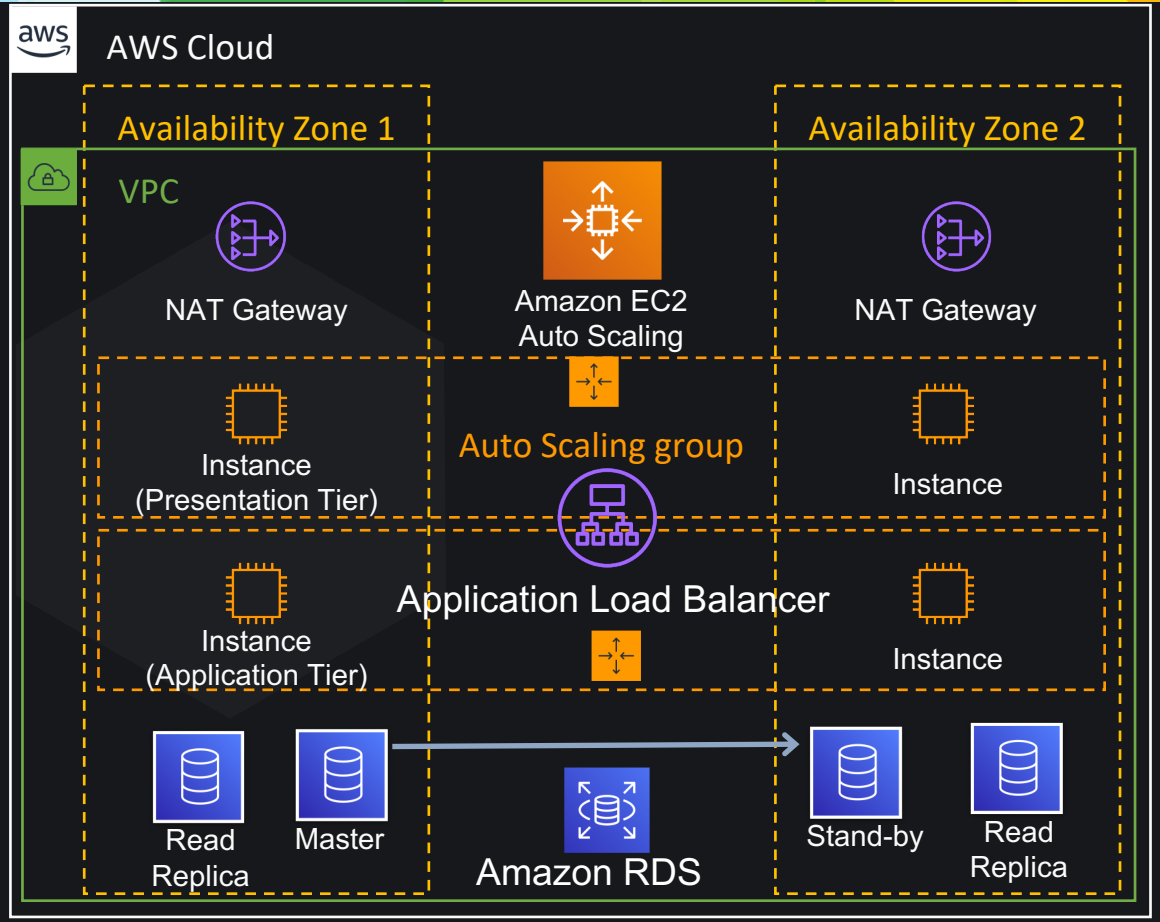

Object Storage File Systems in the Cloud

- Amazon S3 democratized access to scalable cost-effective long-term storage via simple APIs.
- AWS is responsible for capacity planning, storage provisioning, fault-tolerance and long-term durability through replication.
- Could this level of automation be applied to computing as well?

Amazon Simple Storage Service (S3)

Bucket with objects

Enter AWS Lambda

- Execute user-defined stateless functions in response to events on an dynamically managed computing platform (FaaS – Functions as a Service).
- Anatomy of a Lambda function:
  - Coded in a supported programming language (Node.JS, Python, Go, Java, etc.) or BYOR.
  - Up to **3000** parallel invocations executed up to 15 minutes with up to 3008 MB.
  - Scratch workspace of 512 MB (potentially shared across invocations).
  - Pricing in execution blocks of 100 ms with a generous free tier (1M requests and 400.000 GB/s).
  - Triggered in reponse to events (REST API invocation, file upload to S3, etc.)
    - ¡Event-driven computing!

Definition by elimination

- Carriage

- *Horseless* Carriage (1893)

- Wired Phone

- Wireless Phone

Serverless Computing

- Serverless is an architectural pattern that adopts Cloud managed services that feature dynamic resource allocation to allow developers **focus on the application logic**.
- FaaS is an **event-driven**, **pay-per-use** execution model of functions on a computing platform managed by a provider.
- Sometimes used interchangeably, though not everyone agrees:
  - https://www.jeremydaly.com/stop-calling-everything-serverless/

What can Serverless do for science?

- Serverless computing is having a profund impact in how Cloud-native aplications are being developed nowedays …
- … but how can this be applied to **scientific computing**?

Exploiting Thousands of Cores

- PyWren - http://pywren.io/
  - Pywren lets you run your existing python code at massive scale via AWS Lambda
  - Achieves over 40 TFLOPs across thousands of simultaneous cores.
  - Up to 80 GB/sec read and 60 GB/sec write performance to S3.
  - Developed at riselab – Berkeley.

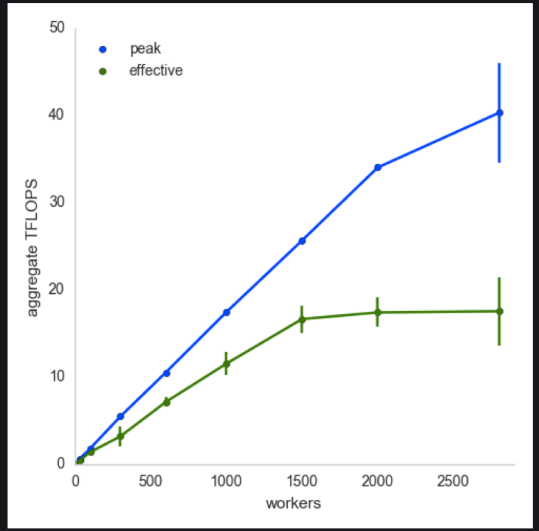

E. Jonas, Q. Pu, S. Venkataraman, I. Stoica, and B. Recht, “Occupy the cloud: distributed computing for the 99%,” in *Proceedings of the 2017 Symposium on Cloud Computing - SoCC ’17*, 2017, pp. 445–451.

Custom Runtime Environments

- SCAR – https://github.com/grycap/scar
  - Highly-parallel event-driven file-processing serverless applications that execute on customized runtime environments provided by Docker containers run on AWS Lambda.
  - Uses UDOCKER to run containers on user space, a development from the INDIGO - DataCloud project.

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.

SCAR's Architecture

- Parallel invocations to Lambda functions that run the user's script in the Docker container to efficiently process data files uploaded to S3 (or invocations to API Gateway)

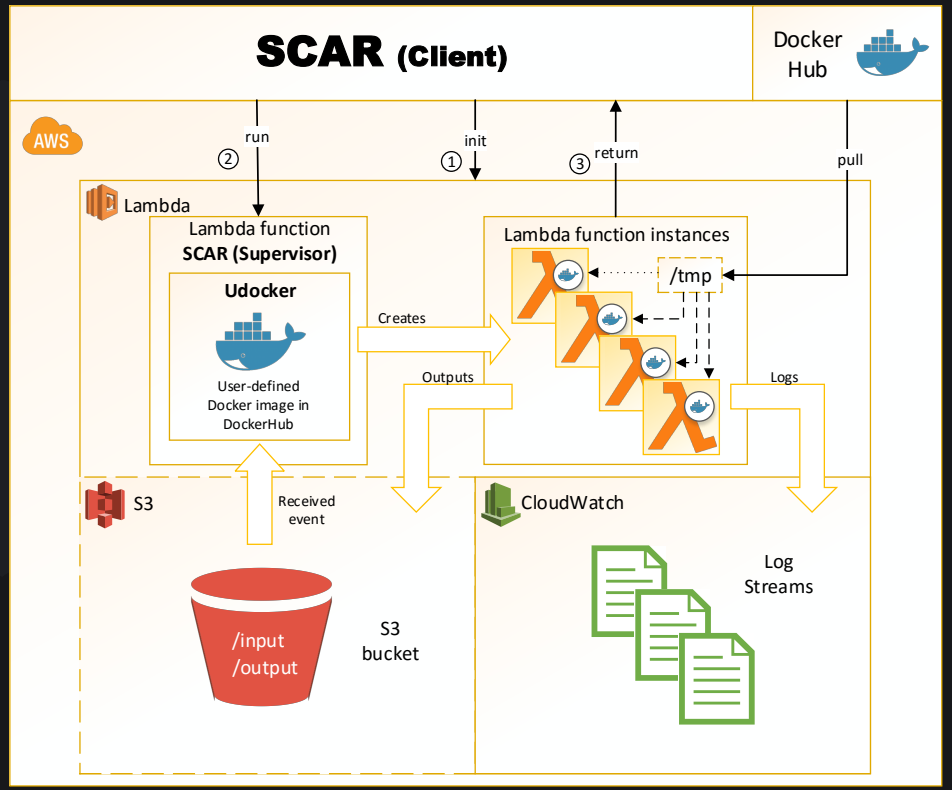

SCAR in the CNCF's Serverless Landscape

- https://landscape.cncf.io/format=serverless&selected=scar

Cloud Native Computing Foundation – Serverless Landscape

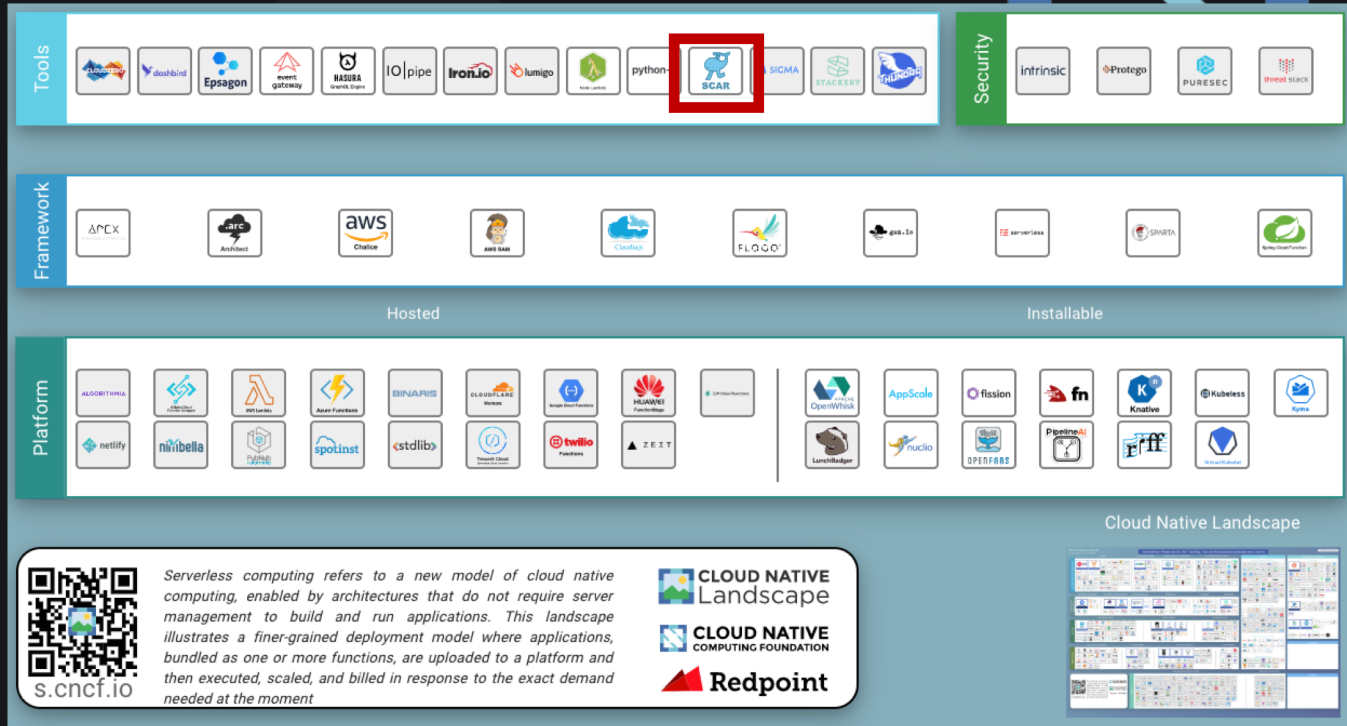

SCAR extension to AWS Batch

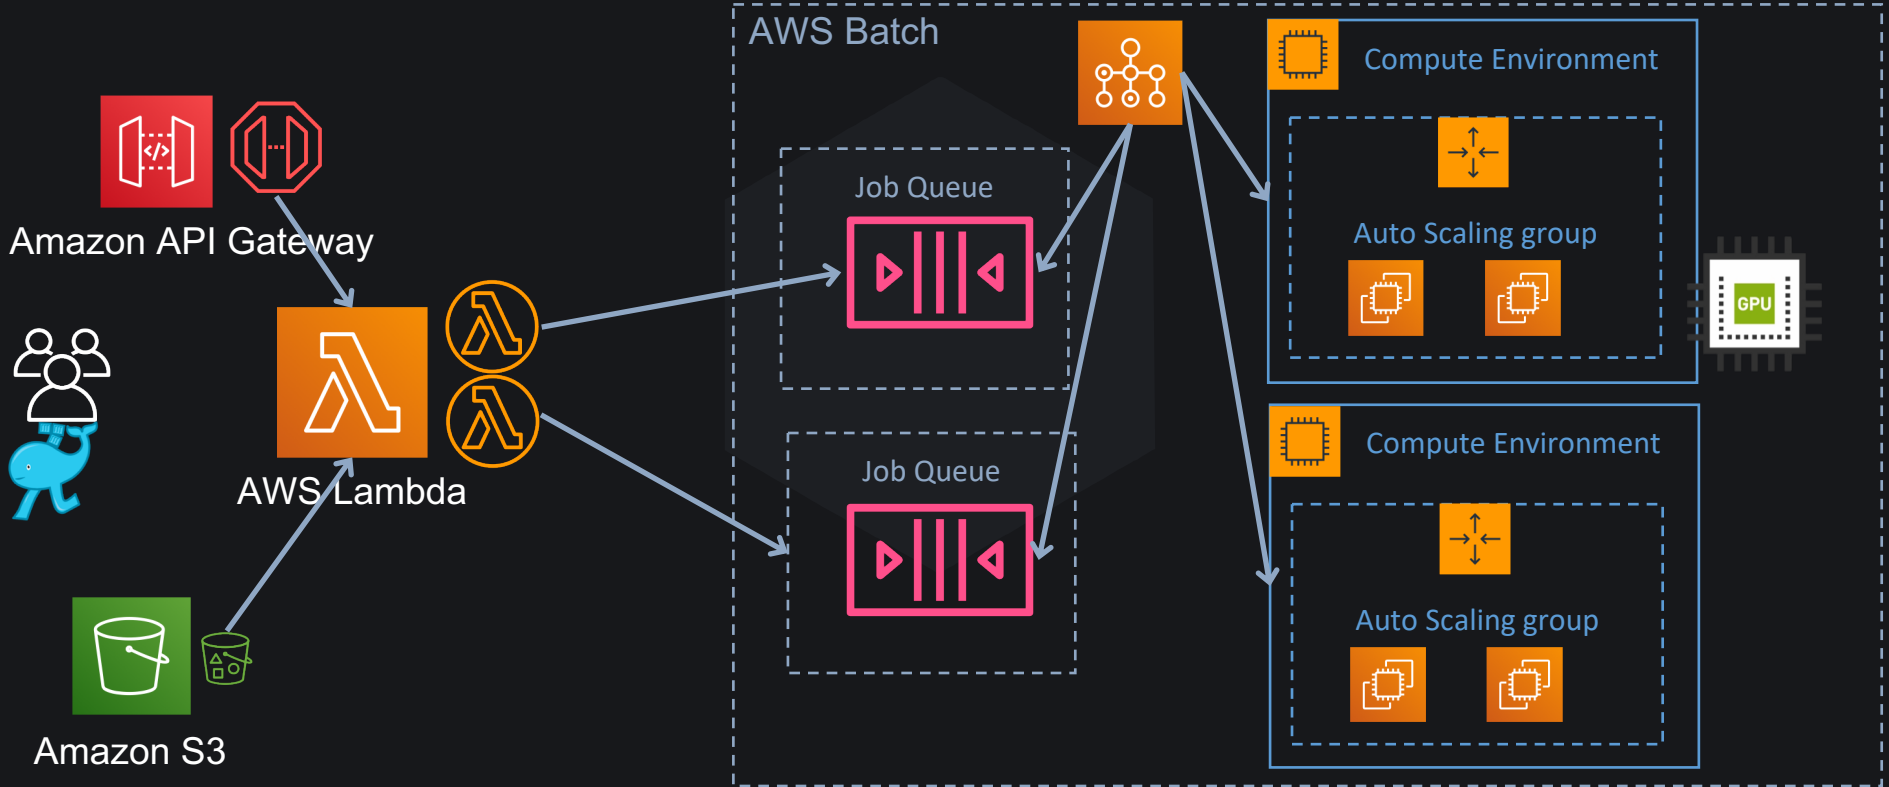

Serverless MapReduce

- MARLA - https://github.com/grycap/marla
  - Deploy a serverless MapReduce processor on AWS Lambda. Files are uploaded to Amazon S3 to trigger the execution of user-supplied Mapper and Reduce functions.
  - Automated data partitioning and parallelism.

V. Giménez-Alventosa, G. Moltó, and M. Caballer, “A framework and a performance assessment for serverless MapReduce on AWS Lambda,” *Futur. Gener. Comput. Syst.*, Mar. 2019.

MARLA Architecture

- The coordinator decides the number of Mappers depending on the dataset size and scalability limits.
- Mappers retrieve a subset of data from S3 in parallel and execute concurrently.

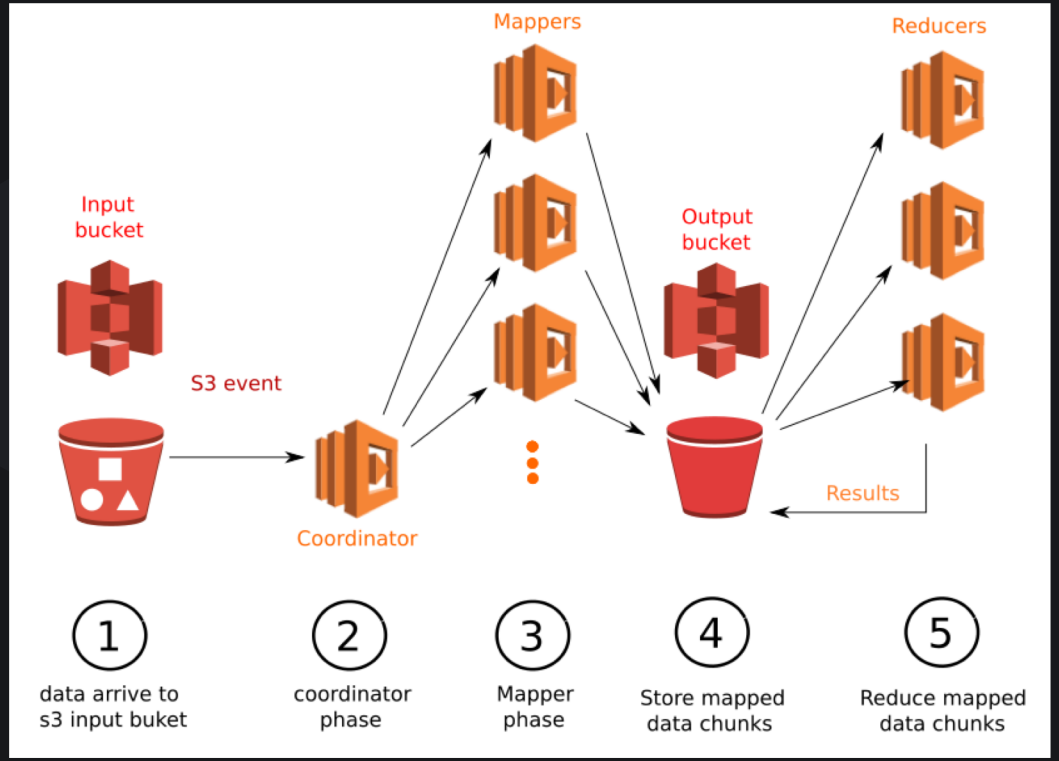

On Performance

- AWS Lambda provides unprecedented levels of elasticity.
- But on a sometimes inhomogeneous platform that may affect coupled executions.

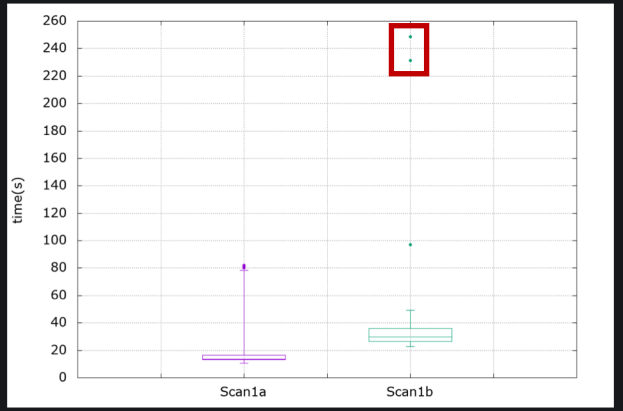

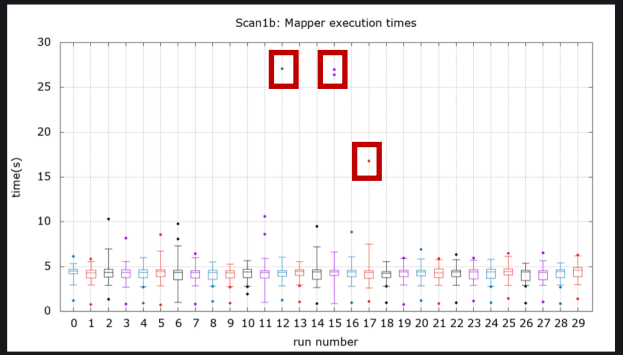

V. Giménez-Alventosa, G. Moltó, and M. Caballer, “A framework and a performance assessment for serverless MapReduce on AWS Lambda,” *Futur. Gener. Comput. Syst.*, Mar. 2019.

Going Serverless for Everyday Computing

- *"Instead of running these tasks on a laptop, or keeping a warm cluster running in the cloud, users might push a button that spawns 10,000 parallel cloud functions to execute a large job in a few seconds from start"*
- GG- https://github.com/StanfordSNR/gg
  - A framework to execute applications on thousands of parallel threads run as Cloud functions (use case of a distributed compiler run on AWS Lambda).

Sadjad Fouladi *et al.*, "From Laptop to Lambda : Outsourcing Everyday Jobs to Thousands of Transient Functional Containers," *USENIX ATC-sbm*, vol. 77, no. 1, 2019.

The Devil is in the ~~Details~~ Costs

- For intensive usage rates, a traditional architecture based on VMs may be more cost-effective.

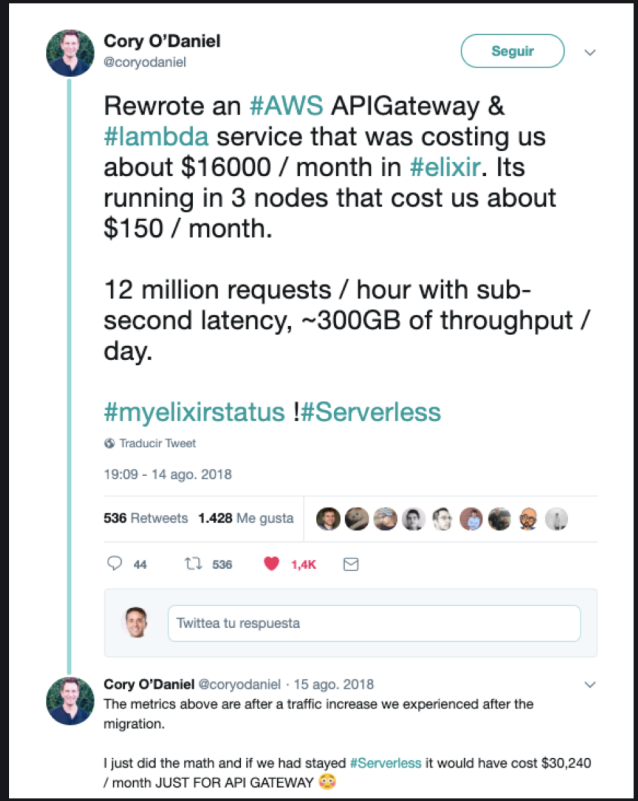

https://twitter.com/coryodaniel/status/1029414666681469952

FaaS: No Silver Bullet

- I’ve got a quick, computationally-intensive task that I need to perform ocassionally in response to a well-defined event that isn’t that sensitive to latency.

Clay Smith - The Ideal FaaS Developer

- Quote Source: https://www.dotconferences.com/2017/04/clay-smith-searching-for-the-server-in-serverless

Serverless … on-premises!???

OPENFAAS

fission

APACHE
OpenWhisk™

riff

f(x)

Anatomy of a FaaS Framework

- Functions packaged as Docker images.
- Gateway to provide REST API to define/invoke functions.
- Monitoring and scaling (at the level of containers, not VMs)

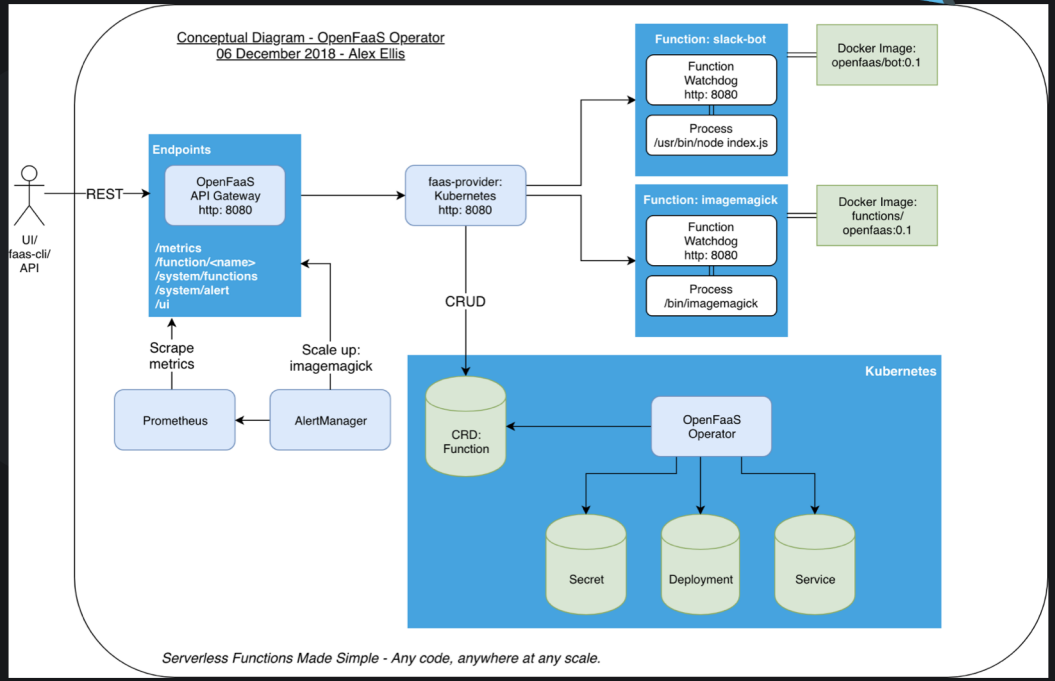

https://docs.openfaas.com/images/of-conceptual-operator.png

OSCAR: Components

- https://github.com/grycap/oscar
- Open-source platform to create **highly-parallel event-driven** file-processing serverless applications that execute **Docker** containers on an **elastic Kubernetes** cluster.
- Partially funded by the EGI Strategic and Innovation Fund.

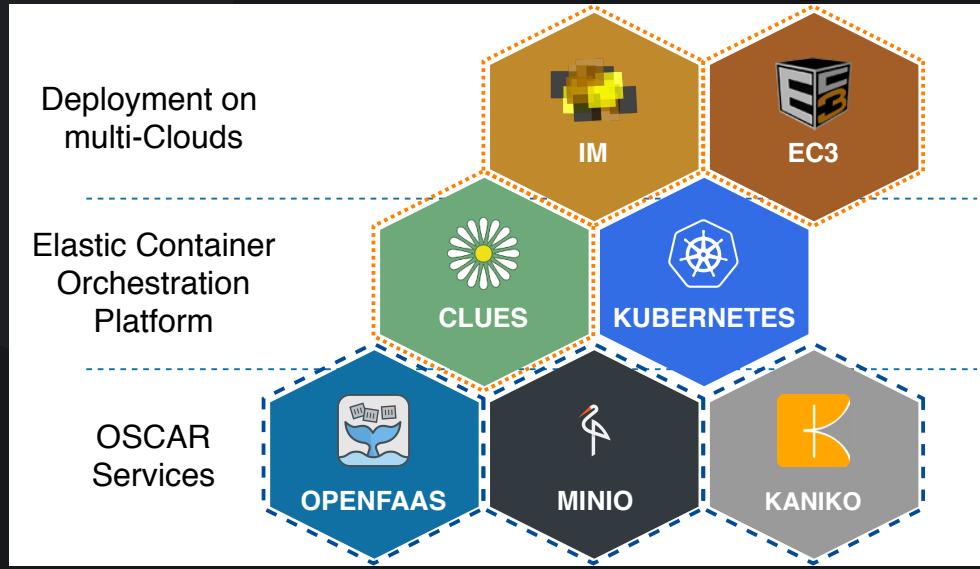

OSCAR Architecture

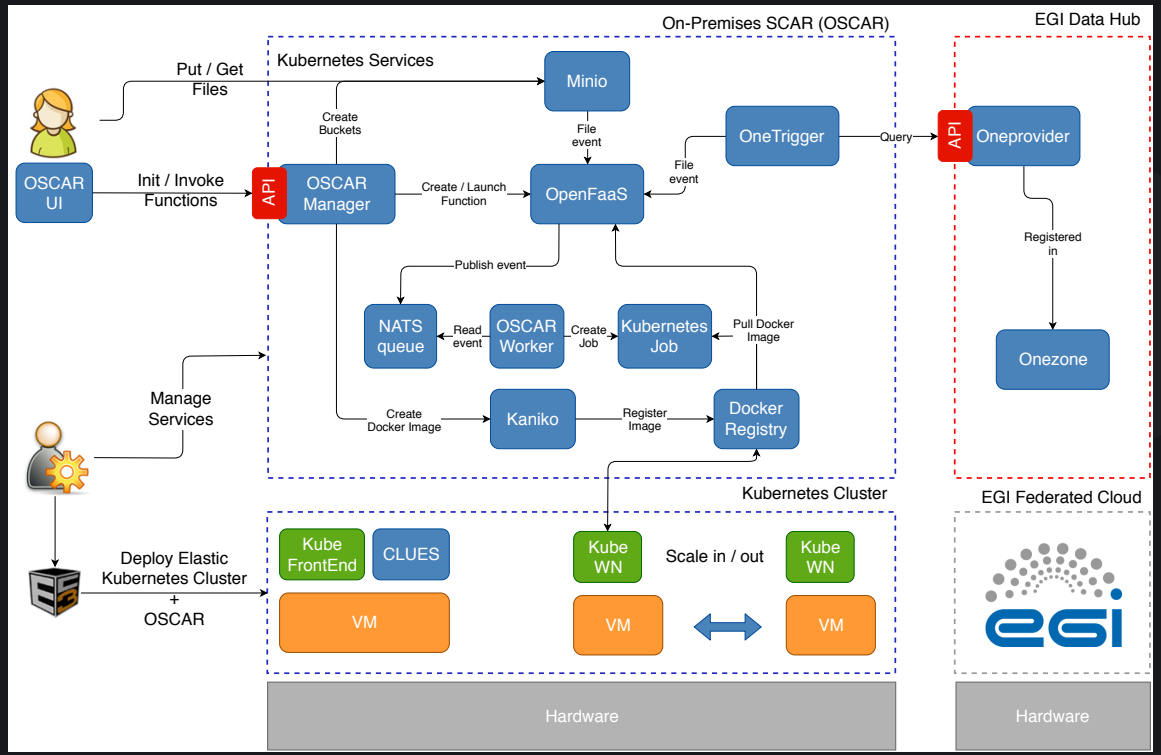

- Users upload file to a storage back-end, which triggers the parallel execution of a user-defined file-processing script ran on a user-defined Docker container.
- Adaptive elasticity of the Kubernetes cluster.

Elastic Kubernetes Cluster

- The Kubernetes cluster dynamically grows and shrinks according to the workload of jobs to be processed.

EKaaS

Elastic Kubernetes as a Service

- EKaaS (Elastic Kubernetes as a Service), funded by the EGI Strategic and Innovation Fund.

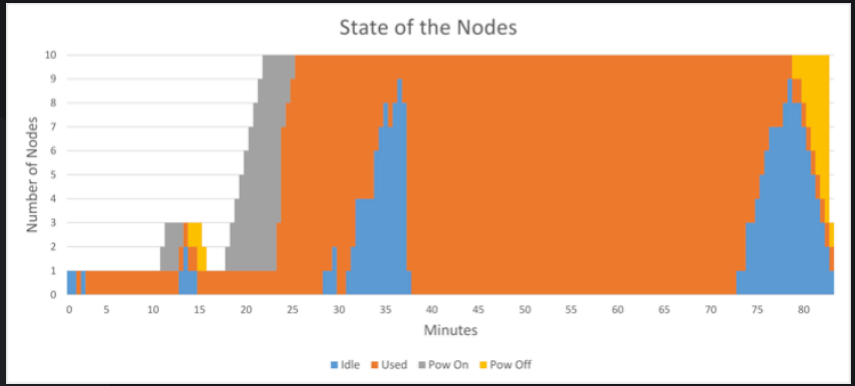

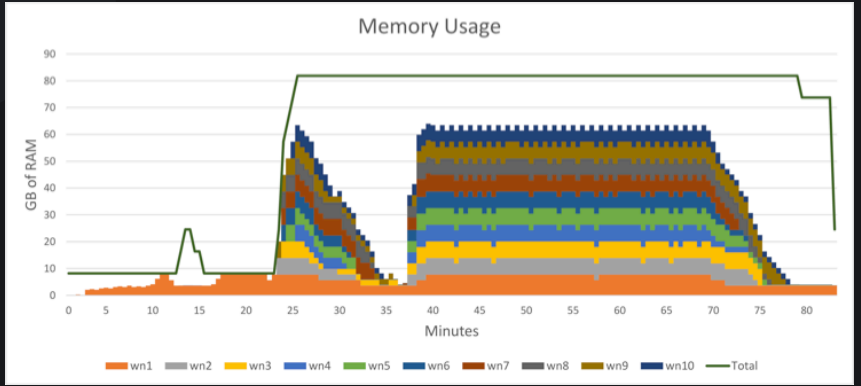

OSCAR: Use Cases

- Plant classification using Deep Learning models trained in the context of deep Hybrid DataCloud.
- Flows of functions to process in parallel video frames on a set of video files.

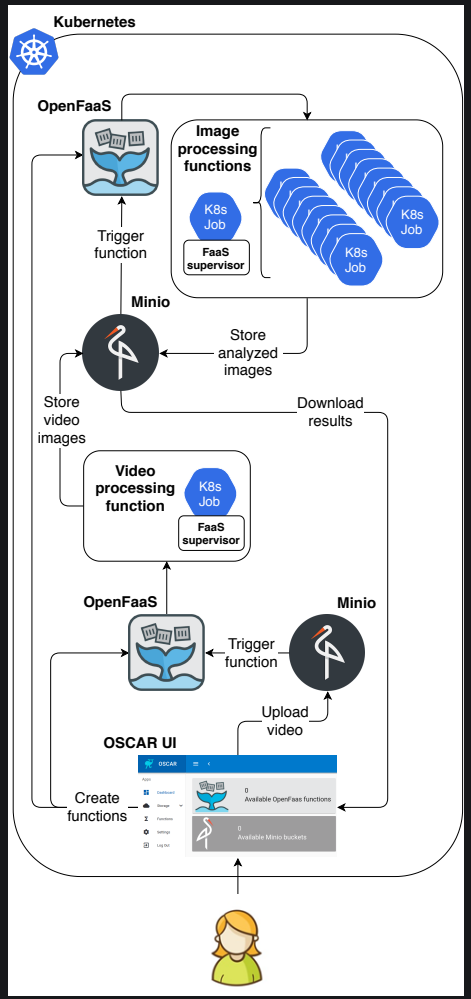

A. Pérez, S. Risco, D. M. Naranjo, M. Caballer, and G. Moltó, “Serverless Computing for Event-Driven Data Processing Applications,” in *2019 IEEE 12th International Conference on Cloud Computing (CLOUD 2019)*, 2019, pp. 414–423.

Integration with GPUs

- OSCAR integrated with virtualized GPU support (rCUDA, etc.).
- Slight overhead due to container start and loading Python libraries.

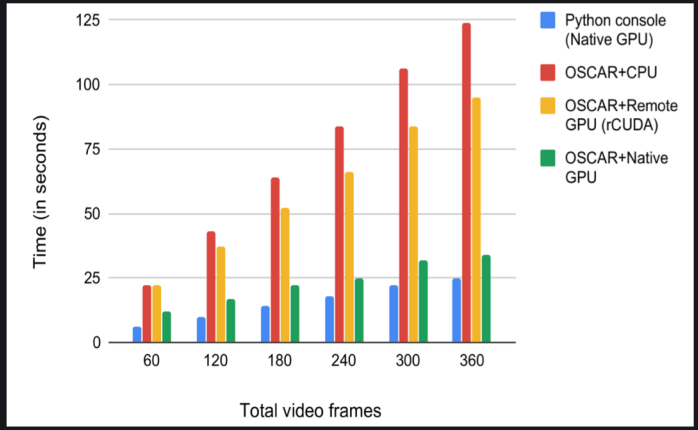

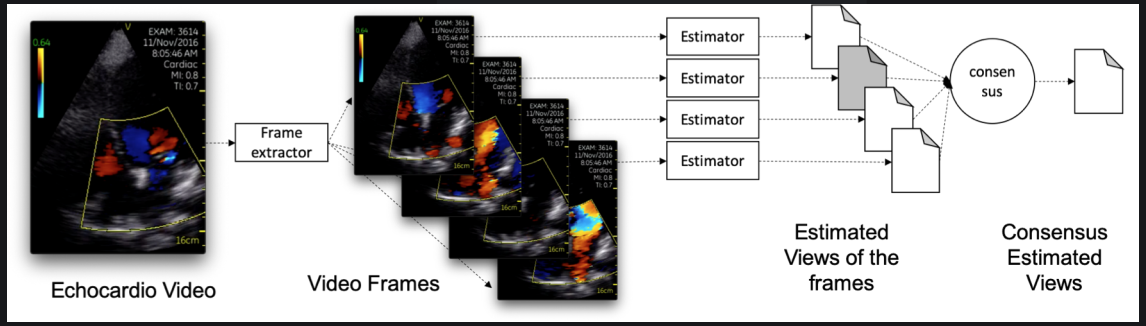

Diana M. Naranjo, Sebastián Risco, Carlos de Alfonso, Alfonso Pérez, Ignacio Blanquer, Germán Moltó, “Accelerated Serverless Computing based on GPU Virtualization,” Journal of Parallel and Distributed Computing. Special issue: Virtualization for Future Computing Systems (*under review*) ,

Conclusions

- Serverless is a computing model to focus on user-level application logic rather than interacting with low-level infrastructure details, typically involving function-based event-driven computing (FaaS).
- Multiple frameworks to support the FaaS computing on-premises managed by a Container Orchestration Platforms (e.g. Kubernetes).

The Future Looks ~~Bright!~~ ~~Cloudless~~

Antonio Goncalves @agoncal · 20 jun.
En respuesta a @HPE
Can't wait for the Computerless Computing
5 2 29

Bruno Borges @brunoborges · 20 jun.
I am actually waiting for the "Inkless Printer, by HP"!
3 1 5

Armen @armdev · 19 jun.
En respuesta a @HPE
hmm, kubernetesless?
4

Corey Quinn @QuinnyPig · 20 jun.
En respuesta a @HPE
Clueless.
16

Clueless

HPE ✓
@HPE

Introducing cloudless computing: A new approach that obliterates the distinction between 'public-ness' and 'private-ness'. Want to know more about what #cloudless computing is? Watch this: hpe.to/6009Eqn0l #HPEDiscover

Traducir Tweet

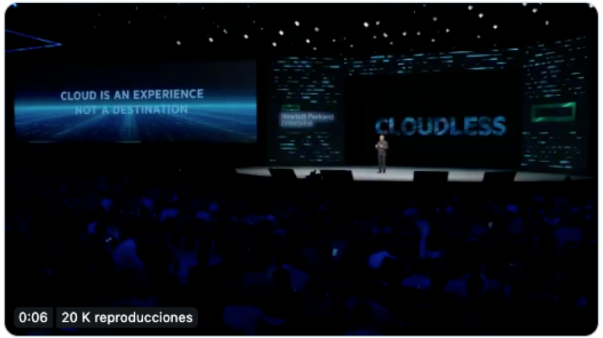

1:09 a. m. · 19 jun. 2019 · SnappyTV.com

10 Retweets **33** Me gusta

https://twitter.com/HPE/status/1141120699169271808

Contact & Acknowledgements

Germán Moltó

Universitat Politècnica de València

gmolto@dsic.upv.es

http://www.grycap.upv.es/gmolto

- SCAR and OSCAR have been partially funded by project BigCLOE (TIN2016-79951-R).
- OSCAR has been partially funded by the EGI Strategic and Innovation Fund.

References

- https://www.computerweekly.com/news/450302728/Australian-2016-census-sabotage-puts-a-question-mark-on-private-cloud
- https://eftm.com/2016/08/census-2016-the-10-million-online-census-what-went-wrong-30681
- https://apo.org.au/sites/default/files/resource-files/2016/11/apo-nid70705-1232016.pdf
- https://www.computerweekly.com/news/450403576/IBM-blamed-for-Australian-census-website-crash
- https://www.zdnet.com/article/australian-2021-digital-census-to-be-built-on-aws/
- https://blog.gigaspaces.com/amazon-found-every-100ms-of-latency-cost-them-1-in-sales/
- https://www.news.com.au/technology/online/hacking/what-does-this-digital-attack-map-tell-us-about-the-alleged-census-attack/news-story/2c06914dec07beca6079801634b99a58
- https://www.huffingtonpost.com.au/2016/08/09/twitter-is-having-a-field-day-over-censusfail_a_21447984/
- https://serverless.com/blog/building-a-better-australian-census-site/
- https://www.jeremydaly.com/stop-calling-everything-serverless/
- https://medium.com/weareservian/getting-started-with-aws-batch-3442446fc62
- https://www.dotconferences.com/2017/04/clay-smith-searching-for-the-server-in-serverless
- https://www.computerworld.com/article/3146568/and-there-she-goes-hpe-jettisons-both-openstack-and-cloud-foundry-initiatives.html

Links to Pictures

- Australia Map: https://emigrara.com/wp-content/uploads/2017/05/Australia-1024x845.jpg
- Kangaroo: https://eacnur.org/blog/wp-content/uploads/2017/07/historia-de-australia_opt-800x400.jpg
- Lake Hillier: http://www.goldfieldsairservices.com/lake-hilliermiddle-island-flight
- Wireless Phones: http://www.actionlinkwireless.com/history-cell-phone/
- Horseless Carriage: https://hackastory.com/vr-storytelling-blog-1-the-horseless-carriage-syndrome/