


JASMIN (STFC/Stephen Kill)

JASMIN and the adoption of cloud-native architecture for managing data and compute at scale

16th Workshop on Meteorological Operational Systems, ECMWF

1 – 3 March 2017

Philip Kershaw[+], Jonathan Churchill[*], Bryan Lawrence[^], Matt Pritchard[+], Matt Pryor[+]

+ NCAS/NCEO Centre for Environmental Data Analysis, RAL Space, STFC;
* Scientific Computing Department, STFC; ^ NCAS University of Reading

JASMIN Introduction

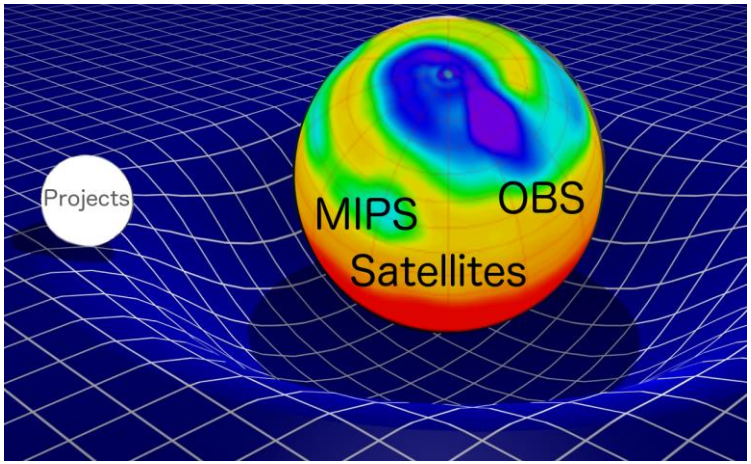

Data gravity associated *with managed data* so that users want to bring their projects to the the JASMIN environment

- JASMIN is a NERC-funded multi-petabyte data analysis facility
  - for the UK environmental science community and their international collaborators.
  - Over 1000 registered users
  - Hosts CEDA data centres supporting 30k users
- A data commons: *bringing the compute to the* [managed] *data* paradigm
  - Managed data – analysis ready: Big Data **Value** 'V'
  - Predominantly climate science and Earth observation
- A response to the challenges of Big Data encountered in this and other research domains.
- In operation since 2012
  - Celebrated 5th birthday earlier this week ☺

Data commons – compute to the data

Make your own VMs, storage and network config

Lotus Batch Compute

Virtualisation

Platform as a Service

Infrastructure as a Service

Software as a Service

Build and run code to run in parallel on shared environment

e.g. OPeNDAP, WMS, IPython Notebook, Web portals

CEDA Archive – Data Cent

Group Workspaces

Access pre-prepared scientific analysis hosts – preconfigured with apps and libraries for community

Provision your own shrink-wrapped VMs e.g. database, analysis, web server

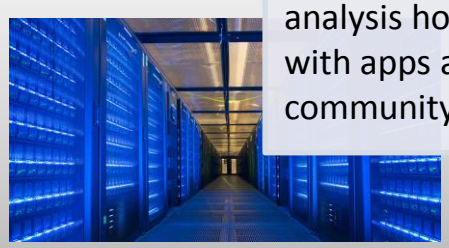

JA

Parallel file system (1 PB); high–speed, non-blocking, low latency networking;

- Batch Compute; Community Cloud; Tape storage

High Level themes

- Why adopt a cloud-native architecture?
  - Elasticity and scaling for Big Data
  - Cloudbursting: migrate into public cloud taking based on based on technical, policy and financial merits

- We already have a cloud architecture for JASMIN what do we need to evolve?
  - *Compute* and *storage* interfaces

Storage: Data-as-a-Service

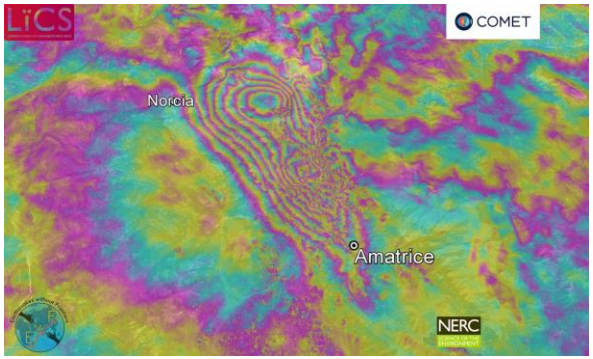

A map of deformation caused by an earthquake
Contact: Tim Wright (University of Leeds)
http://comet.nerc.ac.uk/

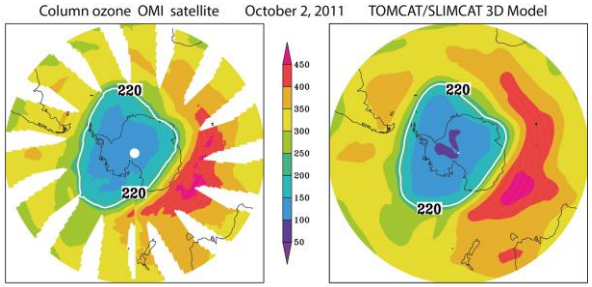

Comparison of satellite coverage vs. model runs
Contact: Martin Chipperfield (Uni Leeds), Wuhu Feng (NCAS), Chris Wilson and Richard Pope (NCEO)
http://www.see.leeds.ac.uk/tomcat

- Growth of data and of the user community supported is driving change in how data access is implemented.
- Ideally, data access should be both performant and ubiquitous for applications consuming them.
- There are two key factors for consideration:
  - the network architecture - enabling performance and isolation
  - the interfaces used to access data.

An e.g. data access scenario: ESA Climate Change Initiative and JASMIN

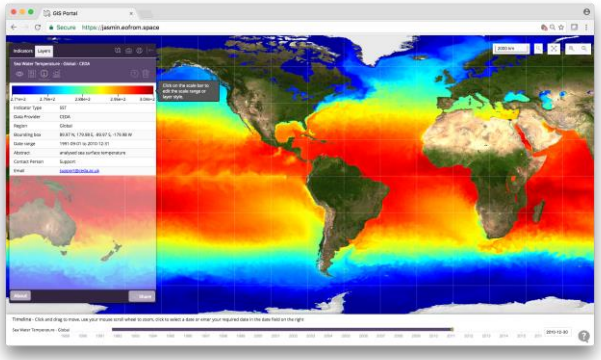

Sea Surface Temperature CCI
Contact: Dr Owen Embury (University of Reading)
GISportal (PML) hosted on JASMIN

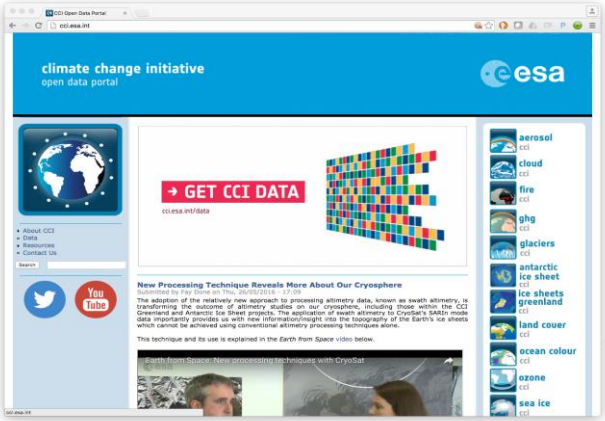

CCI Open Data Portal hosted on JASMIN's cloud

- JASMIN enables a continuous chain:

Data Production:	Curation:	Dissemination:
• ECV (Essential Climate Variable) datasets with **JASMIN Lotus batch compute**	•~180TB on parallel file system • Metadata catalogue records	• Portal and discovery, visualisation, download services hosted on **JASMIN cloud**

- *Lotus* batch compute + parallel FS:
  - Sea Surface Temperature: scientists can generate 30+ years of datasets in just a few days, rather than months or years.
- But . . .

Data access services: Parallel File Systems, POSIX and Cloud

- JASMIN's community cloud allows users to provision virtual machines using an IaaS (Infrastructure as a Service) model

- But there is a fundamental incompatibility between this and parallel file systems at scale
  - Parallel file system: a global uid/gid space under a single administrative authority
  - IaaS model: multiple tenant-defined administrative authorities

- IaaS on JASMIN is segregated into an isolated network
  - Gives full autonomy for tenants
  - access to the data archive and group workspaces via FTP and HTTP interfaces (such as OPeNDAP)

From Parallel File Systems to Object Stores

- Motivations for using a parallel file system in the first place were:
  1) performance for massive data handling, and
  2) ease of management for petascale storage

- Object stores potential to provide a universal interface in the form of the widely adopted S3 REST API:
  - Decouple access policy from POSIX semantics
  - Software-defined solution: can incrementally add new h/w
  - Cost

- Support the main usage classes:
  - batch compute and cloud services from within JASMIN
  - external download
  - Interoperability: migrate easily between on-prem and cloud providers

Object store work with JASMIN

- Proof-of-concept with vendors underway
  - Possibility of dual POSIX/Object store interfaces with some solutions
- HDF (and hence NetCDF4) REST API over object store
  - Development of as part of the European ESIWACE project
- How to address legacy scientific applications and their access to the file system via hierarchical directories?
  - Faceted search systems such as that created for ESGF (Earth System Grid Federation) illustrate alternative approach
  - Mimic directory hierarchy
  - But allow flexible combinations of vocabulary terms to find data
- CEDA File-Based Search project has indexed the whole archive (3.7PB) using ElasticSearch

Compute interfaces and cloud

- Abstraction of compute with cloud
  - APIs allowing the dynamic provisioning of virtual machines
  - But the lack of consistent APIs and formats for sharing virtual machine images has hindered portability between platforms

- Possible solutions for interoperability of cloud compute:
  - Libraries: jclouds, libcloud
  - Terraform
  - Containers technologies

RCUK Cloud Working Group: Pilot project

- Research Councils UK Cloud Working Group
  - goal is to support researchers and technical specialists in the application of cloud computing technologies and services for the research community in the UK
- Exploring a number of areas around cloud adoption:
  - Technical; legal, policy and regulatory; costs
- Pilot project: targeted activity to investigate interoperability across clouds
  - Particle physics chosen as the domain area
  - Use container technologies as a means of abstraction and therefore interoperability across cloud platforms
- Containers allow change from
  - application only to instead
  - an encapsulation of application + dependency bundle

RCUK Cloud Working Group
Pilot: federation across Google and Azure

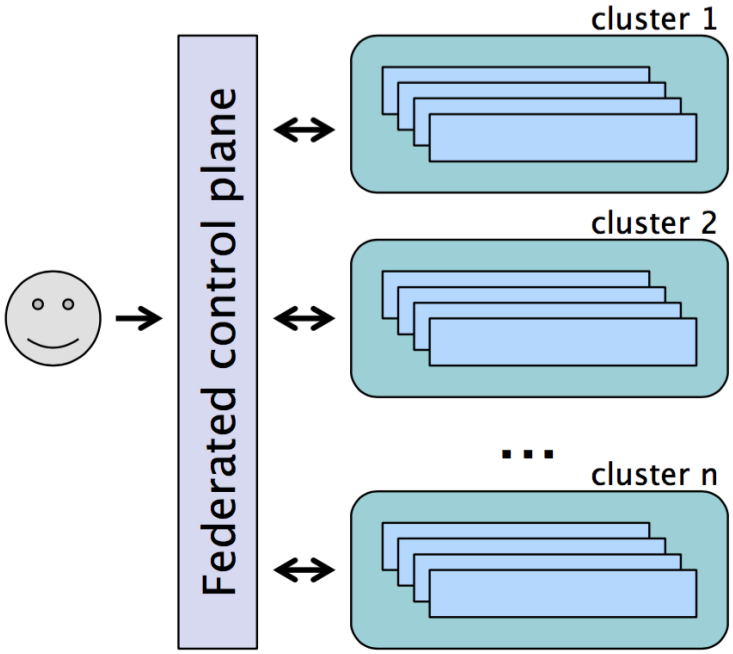


```
$ kubectl --context=federation-cluster get clusters
NAME                 STATUS    AGE
azure-west-europe    Ready     6h
gce-europe-west1     Ready     13h
gce-us-central1      Ready     5m
gce-us-east1         Ready     13h
gce-us-west1         Ready     5m
```

Example with 1 cluster in Azure & 4 clusters in different regions in Google Cloud Platform
- *single API to deploy applications on both Azure and Google Cloud Platform*
- *run one command to create squids in all 5 clusters*

- Courtesy of Andrew Lahiff, STFC

kubernetes

Containers and JASMIN: work to date and plans for the future

Attendees at ESA Summer school, ESRIN used OPTIRAD environment – Credit ESA

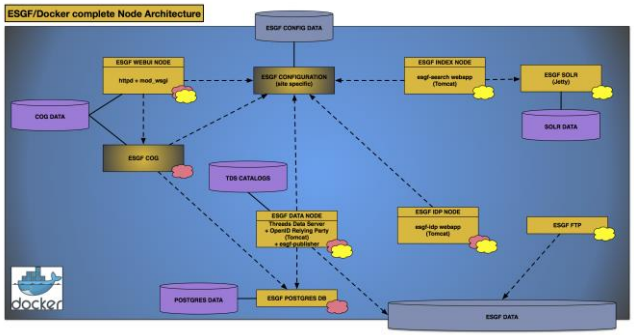

https://github.com/ESGF/esgf-docker/wiki

- OPTIRAD
  - ESA-funded project to collaborative research environment for land-surface data assimilation
  - JupyterHub and IPython.parallel
  - Used to provide training environment at ESA summer school
  - Docker containers with *Swarm* orchestration
- Build on OPTIRAD and make a generic Jupyter (IPython) Notebook service for JASMIN
  - Replace Swarm with *Kubernetes* for orchestration
  - Fully automated deployment with Ansible playbook(s)
  - Looking at integration with MetOffice Informatics Lab work with Jade – using Dask scheduler
- Kubernetes
  - *Pod* encapsulates an application container(s), storage resources, network IP, and container run options
- Adopt container-based solutions for ESGF and Copernicus services

Conclusions

- JASMIN: data commons – bring compute to the (managed) data
- Evolution of Data-as-a-Service is necessary for scale-up for demands of Big Data and the needs of a multi-tenancy hosted computing environment
- Careful stepped approach for object store implementation is required
  - from deployment to full adoption for user community
- There are new possibilities for abstraction of compute and storage
  - Developments in container technology – Kubernetes
  - Possible future convergence with batch compute / HPC with *Shifter*
- Docker, Kubernetes + S3 API for object stores

  → Complete solution for migration of applications between on-prem private and public cloud

Further Information

- CEDA and JASMIN:
  - http://www.jasmin.ac.uk/
  - http://www.ceda.ac.uk/
- JASMIN paper

  Lawrence, B.N. , V.L. Bennett, J. Churchill, M. Juckes, P. Kershaw, S. Pascoe, S. Pepler, M. Pritchard, and A. Stephens. **Storing and manipulating environmental big data with JASMIN.** *Proceedings of IEEE Big Data 2013, p68-75,* doi:10.1109/BigData.2013.6691556
- ESA Climate Change Initiative Open Data Portal
  - http://cci.esa.int/
- ESNet Science DMZ
  - http://fasterdata.es.net/
- CEDA ESGF node
  - https://esgf-index1.ceda.ac.uk/projects/esgf-ceda/
- ESGF ICMWG (International Climate Network Working Group)
  - http://icnwg.es.net/
- Research Councils UK Cloud Working Group
  - https://cloud.ac.uk/
- philip.kershaw@stfc.ac.uk, @PhilipJKershaw