



JASMIN (STFC/Stephen Kill)

### The JASMIN Data Analysis Facility for the Environmental Sciences Community and the Role of Data-as-a-Service

### Session International Applications: Sharing Environmental Big Data, Part I AMS 22-26 January 2017, Seattle

**Philip Kershaw**<sup>+</sup>, Jonathan Churchill<sup>\*</sup>, Bryan Lawrence<sup>^+</sup>, Matt Pritchard<sup>+</sup>, Matt Pryor<sup>+</sup> (+ JASMIN and CEDA teams)

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





**National Centre for** 

Earth Observation

NATURAL ENVIRONMENT RESEARCH COUNCIL

# **JASMIN** Introduction



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

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







ATURAL ENVIRONMENT RESEARCH COUNCIL









ATURAL ENVIRONMENT RESEARCH COUNCIL





European contribution to HiresMIP alone is expected to exceed 2 PB







Sentinel missions data rate: ~6PB/year



National Centre for Atmospheric Science ATURAL ENVIRONMENT RESEARCH COUNCIL



National Centre for Earth Observation NATURAL ENVIRONMENT RESEARCH COUNCIL

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



Centre for Environmental Data Analysis

### An e.g. of data access scenarios: ESA Climate Change Initiative and JASMIN



Sea Surface Temperature CCI Contact: Dr Owen Embury (University of Reading)



CCI Open Data Portal hosted on JASMIN's cloud

National Centre for Atmospheric Science TURAL ENVIRONMENT RESEARCH COUNCIL



**National Centre for** Earth Observation NATURAL ENVIRONMENT RESEARCH COUNCIL

- JASMIN enables a continuous chain from:
  - ECV (Essential Climate Variable) production – with Lotus batch compute
  - Curation: ~180TB on parallel FS
  - External dissemination: Portal and download services hosted on JASMIN cloud
- *Lotus* batch compute:
  - Sea Surface Temperature: scientists can generate 30+ years of datasets in just a few days, rather than months or years.



Centre for Environmental Data Analysis



### Data-as-a-Service: Example 1 – Data Transfer Zone

- Dedicated *Data Transfer Zone* (DTZ) outside the institutional firewall ۲ for WAN data movement.
- This is based on the ESnet concept of a "Science DMZ" ullet
- Experience with CMIP5 and other large projects has demonstrated the need and value
- The DTZ concept is now being extended for download services for ۲ all classes of user.
- This pattern is being piloted for the deployment of ESGF (Earth System Grid Federation) software,
  - Data download services hosted in the DT7
  - Web portals and web services deployed in the private cloud environment where they can be more easily administered.







Centre for Environmental Data Analysis



### Application hosting on JASMIN Community Cloud



Forestry TEP and Polar TEPs hosted on JASMIN-CEMS



CCI Open Data Portal



Majic interface to Jules Land-surface model on JASMIN



#### NERC Environmental Workbench



Centre for Environmental Data Analysis

SCIENCE AND TECHNOLOGY FACILITIES COUNCIL NATURAL ENVIRONMENT RESEARCH COUNCIL



EOS Cloud – Desktop-as-a-Service for Environmental Genomics







OPTIRAD – JupyterHub IPython Notebook hosting



### DaaS Example 2: 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
  - full autonomy for tenants
  - access to the data archive and group workspaces via FTP and HTTP interfaces (such as OPeNDAP)









Mellanox Connect-X4 Dual port 100Gb QSFP+ DA Dell R730XD servers. VXLAN/NV|GRE and Erasure Coding offload in h/w Mellanox Dual MSN2100 16 port x 100G switch/routers



science & Technology Facilities Council Rutherford Appleton Laboratory





- Hardware
  - 100G networking
- Software parallelisation:
  - Traditional load balancing and/or
  - Container-based alternative / combination
  - Build on experience with Swarm and Kubernetes on other projects





National Centre for Earth Observation





### 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.
  - → These are requirements for any successor technology
- Object store potential benefits:
  - Universal interface for access whether inside or outside JASMIN, Cloud service model or other
  - Software-defined solution: can incrementally add new h/w
  - Cost
  - interoperability with public clouds
- Object store work with JASMIN:
  - Proof-of-concept with vendors underway
  - Development of a HDF (and hence NetCDF4) server system with REST API that can be deployed over object stores as part of the European ESIWACE project









### POSIX Applications and Object Stores

- 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)

- Built on Apache Solr



- The ESGF DRS (Data Reference Syntax) defines a set of vocabulary terms indexed from datasets which together uniquely describe it.
  - Mimic directory hierarchy
  - But allow flexible combinations of vocabulary terms to find data
- File-based Search project at CEDA
  - ElasticSearch 😓 elastic
  - Indexing the data archive (3.7 PB)





National Centre for Earth Observation





- JASMIN: data commons bring compute to the (managed) data
- Key data flows and service models
  - Bulk data transfer
  - Cloud and POSIX
  - \_\_\_\_

science & Technology Facilities Council Rutherford Appleton Laboratory

- Require fundamental thinking from bottom up to meet application use cases:
  - network architecture
  - Interfaces to applications
- Careful stepped approach for implementation is required
  - from deployment to full adoption for user community
- 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.







ATURAL ENVIRONMENT RESEARCH COUNCIL



# **Further Information**

- CEDA and JASMIN:
  - <u>http://www.jasmin.ac.uk/</u>
  - <u>http://www.ceda.ac.uk/</u>
- 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
  - <u>http://cci.esa.int/</u>
- ESNet Science DMZ
  - <u>http://fasterdata.es.net/</u>
- CEDA ESGF node
  - <u>https://esgf-index1.ceda.ac.uk/projects/esgf-ceda/</u>
- ESGF ICMWG (International Climate Network Working Group)
  - http://icnwg.es.net/
- philip.kershaw@stfc.ac.uk, @PhilipJKershaw







NATURAL ENVIRONMENT RESEARCH COUNCIL