European Weather Cloud – Introduction

Online Computing Training Week
3 November 2023

Roberto Cuccu (ECMWF)

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

European Weather Cloud

- Pilot project started 2019 by ECMWF and EUMETSAT
- ECMWF new operational infrastructure in ECMWF Bologna Data Centre
- EUMETSAT running on public cloud infrastructure
- **Start of Operations on 26 September 2023**

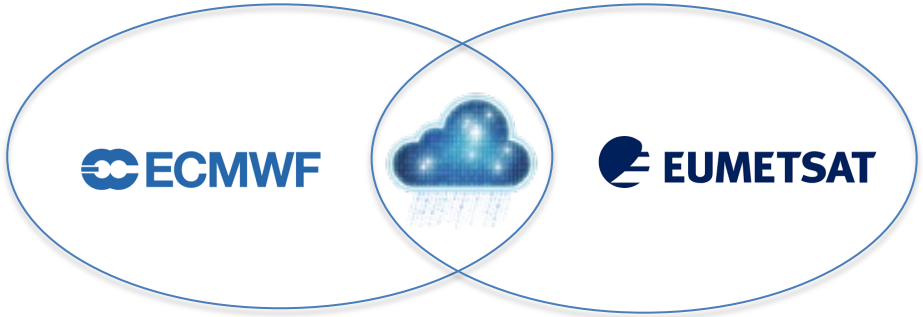

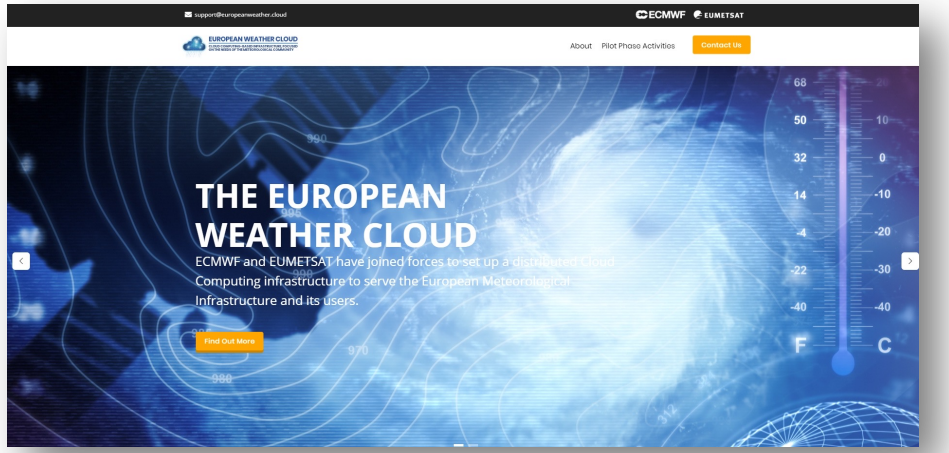

www.europeanweather.cloud

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Who is it for?

Member and Cooperating States

Research & Development

ECMWF Special Projects
EUMETSAT annual R&D calls

EMI Partners (e.g. EUMETNET)

Internal use at ECMWF and EUMETSAT

Member and Cooperating States usage authorized by Computing Representatives

- Access requests via Computing Representatives or EWC support

ECMWF *EUMETSAT*

EWC provides compute capacity and access to ECMWF and EUMETSAT data holdings, additional external data sets and allows users to easily share their own data with others.

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Benefits to the users

- Online access to the cloud computing resources
- Flexibility in provisioning, managing and deleting resources on-demand
- Data availability and data locality for processing
- Community : knowledge, applications, synergies, collaborations

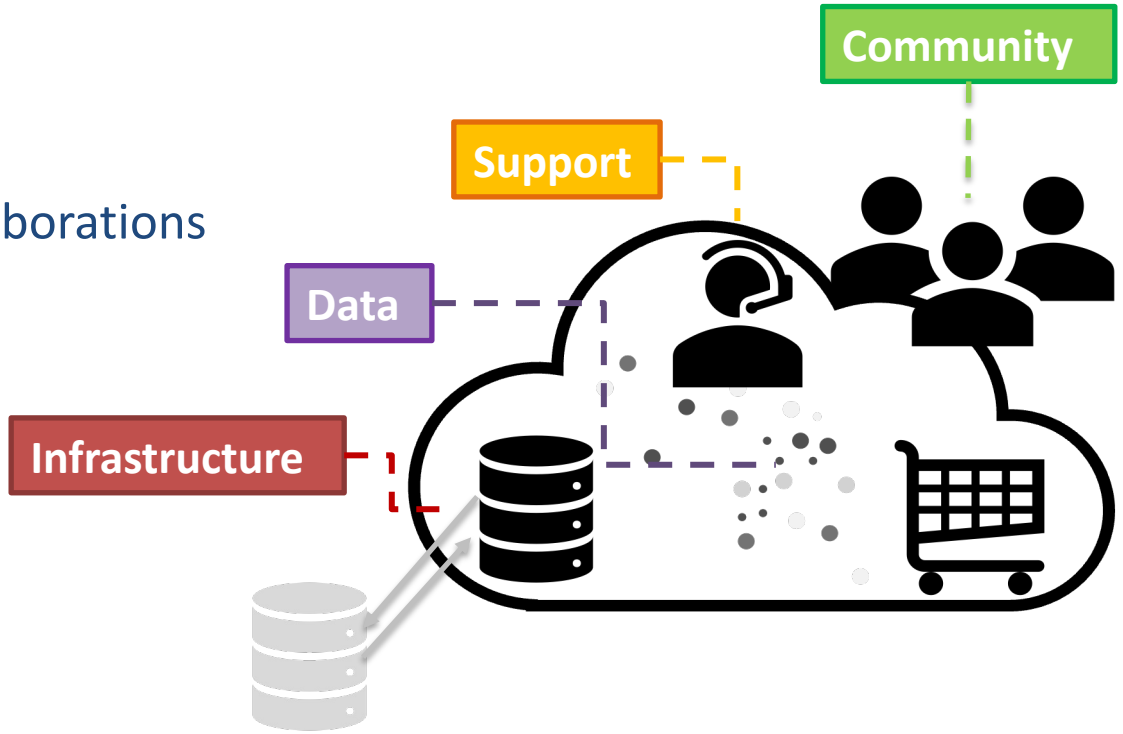

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Cloud Service Description

The service consists of cloud computing resources provided by ECMWF and EUMETSAT and controlled by cloud management software Morpheus.

Virtual Environments

- Isolated cloud tenancy
- Virtual Machines
- Block and Object Storage
- Virtual Networking

. . . . .

Resources Management

- Blueprints and Instance types
- Automation tools
- Monitoring and Reporting
- Metering and Accounting Service

. . . . .

Support and Collaboration

- Support Portal
- Knowledge Base Documentation
- Discussion Platform

. . . . .

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Data Access from EWC

Combined set of “pull” and “push” data access services:

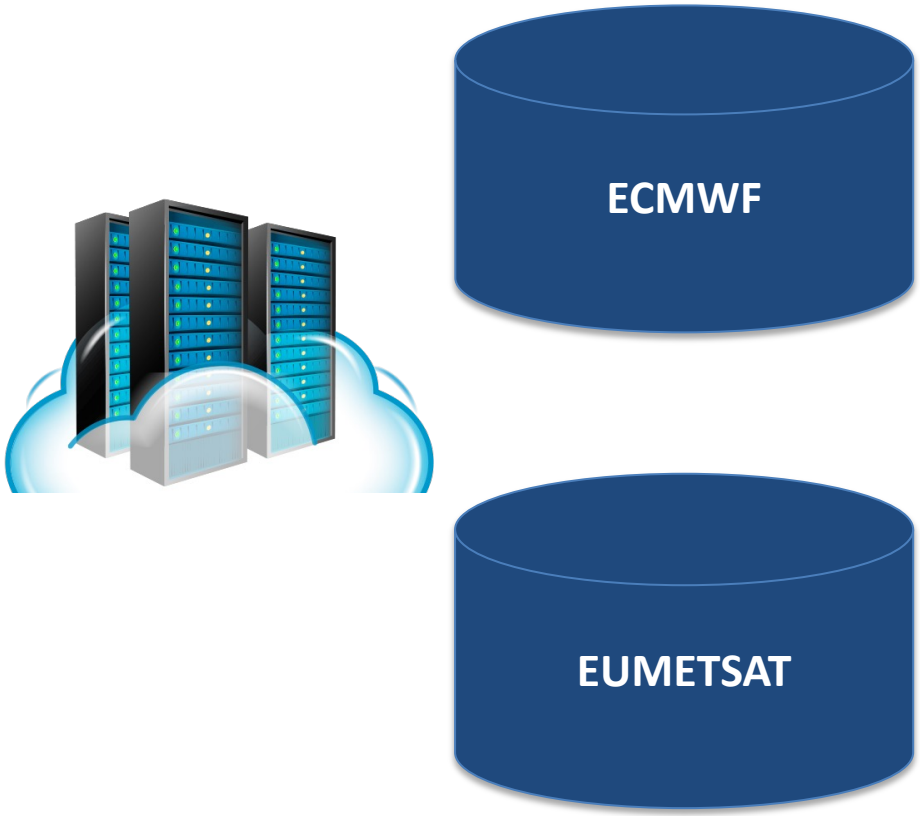

- Meteorological Archival and Retrieval System (MARS)
  - ECMWF Petabytes-scale data archive providing APIs for data discovery and retrieval
- ECMWF Production Data Store (ECPDS)
  - Data dissemination service for customised data delivery
- Copernicus Climate and Atmospheric Data Stores (CDS/ADS)
  - Copernicus Climate Change (C3S) and Atmospheric Monitoring (CAMS) services data
- EUMETSAT Data Store & Data Tailor
  - Access to all EUMETSAT meteorological, climate and ocean data through a suite of APIs, and incorporating data tailoring capability
- EUMETCast Terrestrial
  - Near-real time data delivery via terrestrial network

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

High Level Design

www.europeanweather.cloud

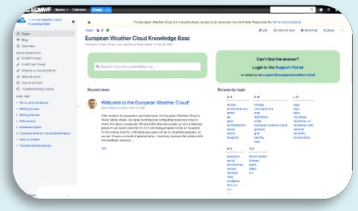

Knowledge Base

Support Portal

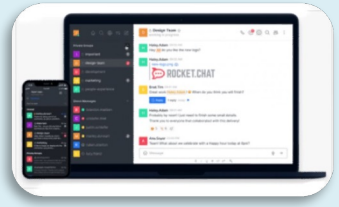

Discussion Platform

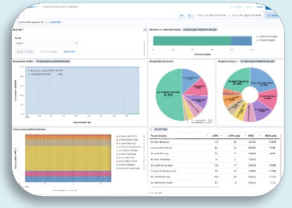

Accounting Platform

User Space

ECMWF MORPHEUS

EUMETSAT MORPHEUS

Cloud Orchestrator

ECMWF openstack. ceph

EUMETSAT openstack. ceph

Cloud and Storage Infrastructure

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

ECMWF Operational infrastructure

- New Cloud Infrastructure ready in ECMWF's Bologna Home
  - Collocated with other key ECMWF Computing and Data Services
- 2 Production clouds - one on each computer hall
  - Based on Openstack and Ceph

Cores	5632
Memory	53 TB
Storage	4.2 PB usable
GPUs	32 x A100 80 GB

- Resources redistributed to the Member and Cooperating States
- Resources allocations managed by the Computing Representative

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Getting a tenancy workflow

Default flow

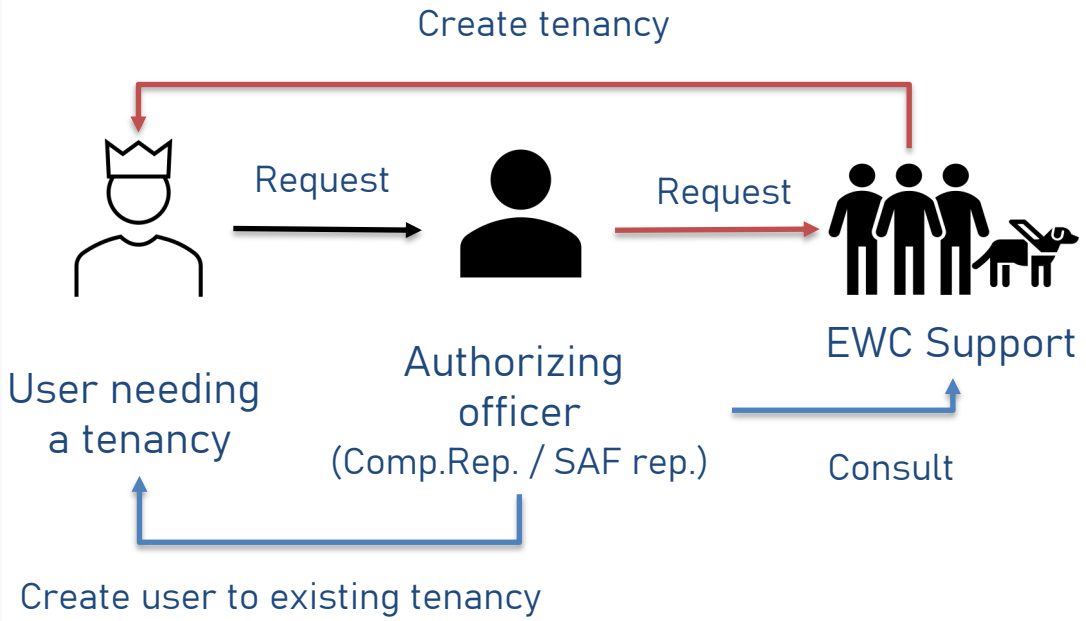

Tenancy will be created to either ECMWF or EUMETSAT side based on the primarily use case

- Tenancy can have resources at both sides
- Member states can have tenancies at both sides

Alternative flow

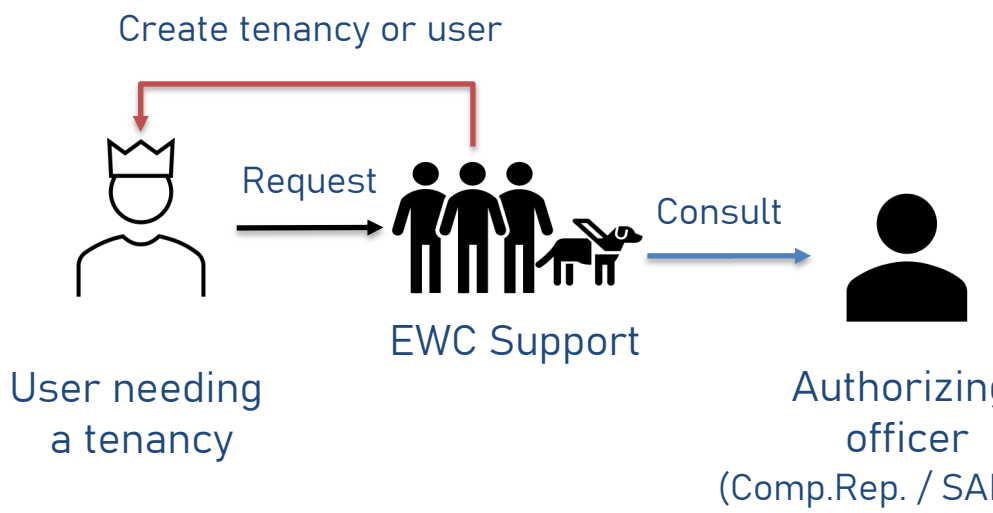

Users may be directly in contact with EWC support in which case support will consult authorizing officer or forward the request

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

ECMWF Special Projects and EUMETSAT R&D Call

ECMWF Special Projects can also include EWC resources in their application, closes each year on 30 June

- The scope includes experiments or investigations of a scientific or technical nature, undertaken by one or more Member States, likely to be of interest to the general scientific community
- ”Late request” possible after deadline

Annual **EUMETSAT Research & Development call** closes each year on 30 June

- Objectives on improving, development and using products in applications and using the cloud infrastructure
- Fast-track projects available anytime of the year for small projects

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

EWC Access – Morpheus Platform

Morpheus: the cloud management platform capable of orchestrating many clouds. Each member of the federation offers a Morpheus instance, to access the federated cloud infrastructures.

Tenants: each project with access to the European Weather Cloud is called a Tenant.

Clouds: Every partner in the European Weather Cloud Federation makes available their infrastructure to the users. Each one of them is seen as a Cloud in Morpheus.

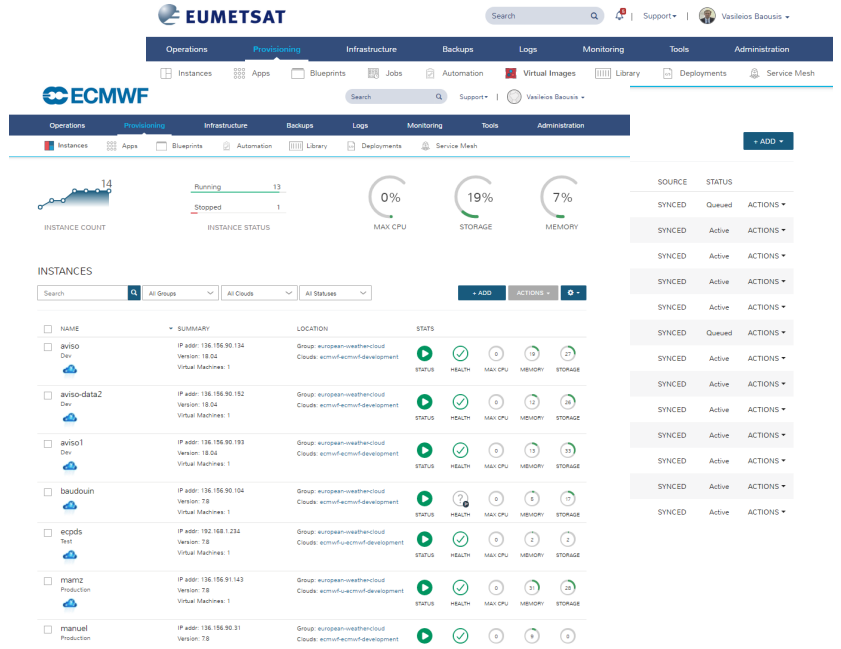

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

ECMWF

Search | Support | Roberto Cuccu

Operations | Provisioning | Library | Infrastructure | Backups | Monitoring | Tools | Administration

Dashboard | Reports | Analytics | Guidance | Costing | Approvals | Activity

INSTANCE STATUS

15
instances

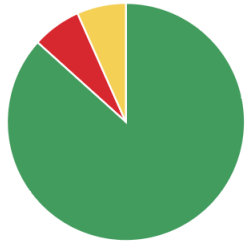

INSTANCES BY CLOUD

3
clouds

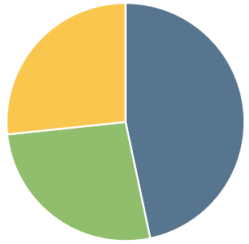

DAILY CLOUD INSTANCES

Week | Month

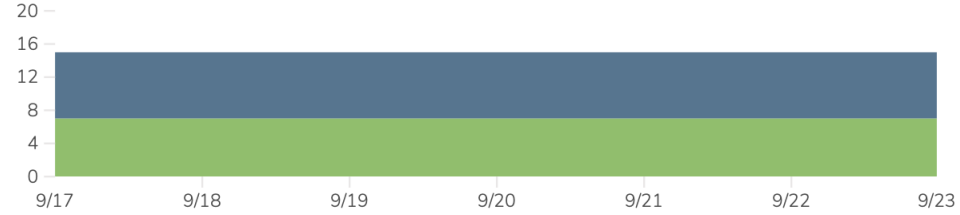

GROUP WORKLOADS

CLOUD WORKLOADS

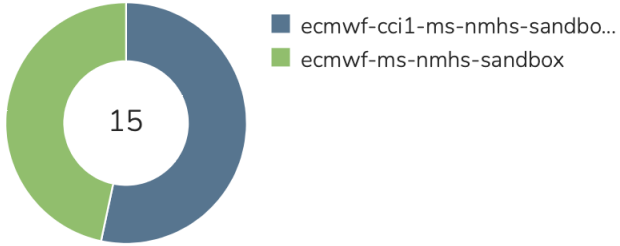

CLUSTER WORKLOADS

NO DATA

LOG HISTORY

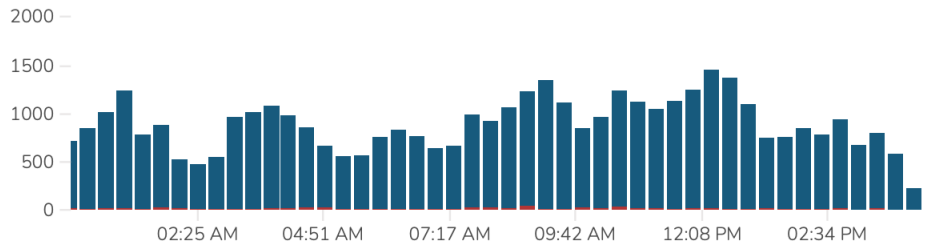

LOG TRENDS

All | Errors | Warnings

Received disconnect from port Bye Bye preauth

pod_workers.go Error syncing pod skipping err failed to StartContainer for openstack-...

I0923 scope.go RemoveContainer containerID

Disconnected from authenticating user root port preauth

I0922 scope.go RemoveContainer containerID

Virtual Machines provisioning

Provisioning-> Instances -> +ADD

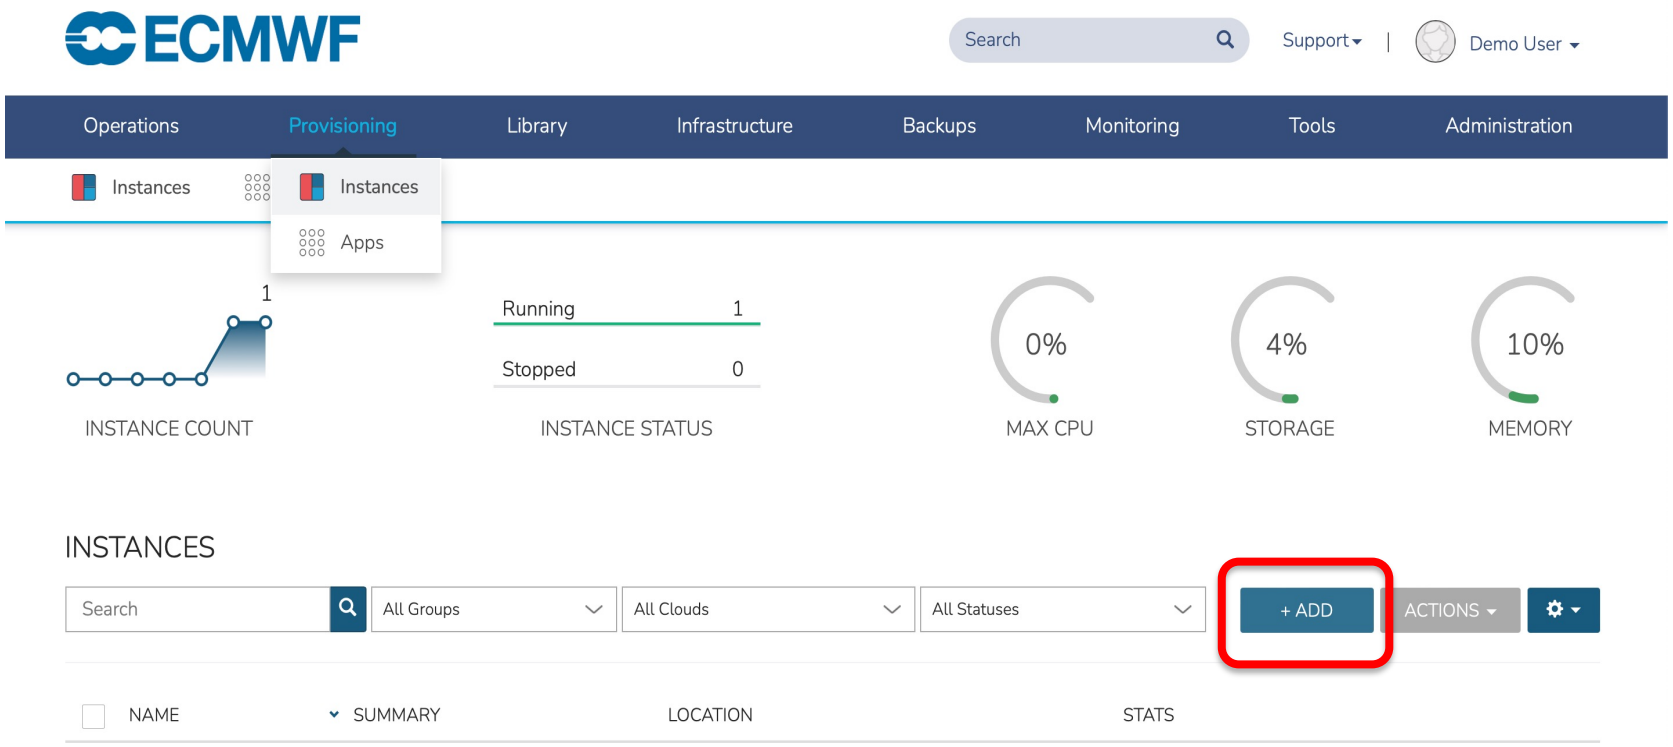

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Virtual Machines provisioning

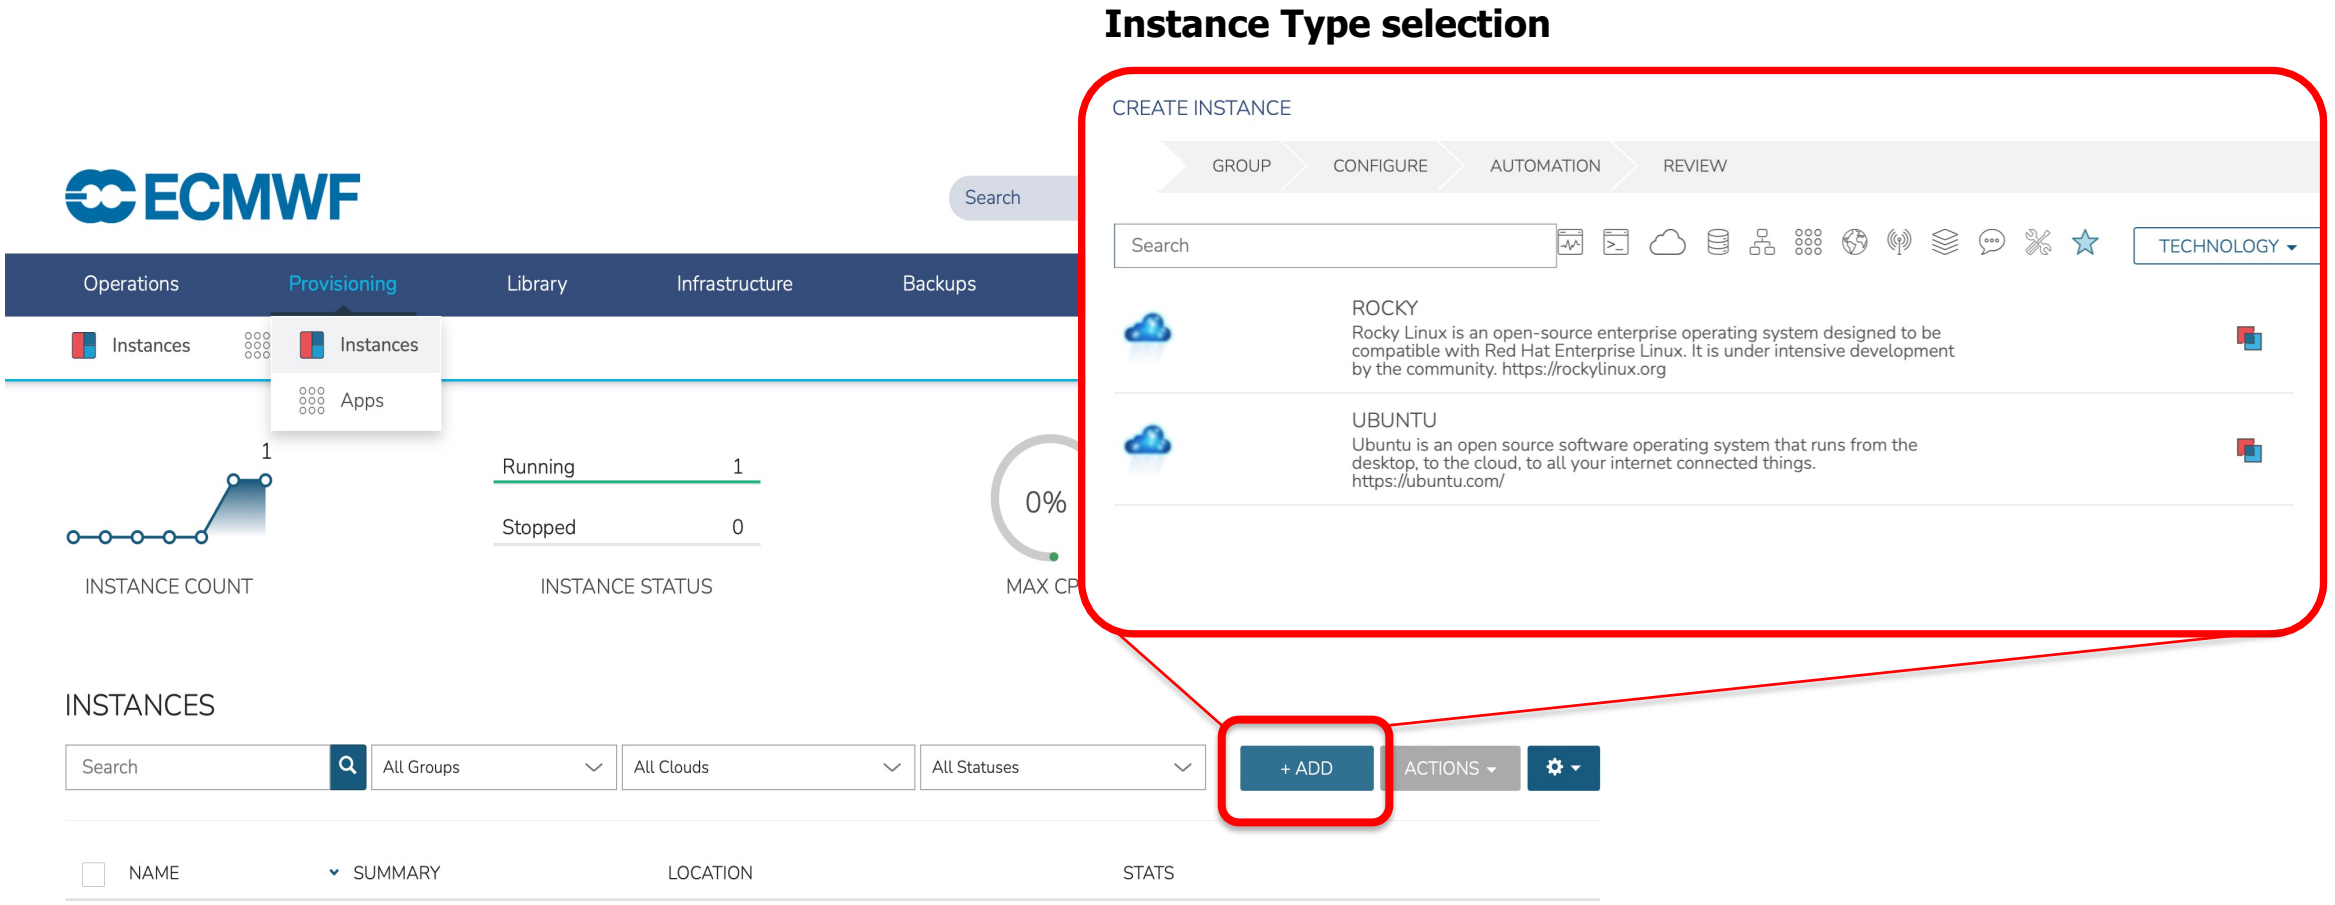

Instance Type selection

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Virtual Machines provisioning

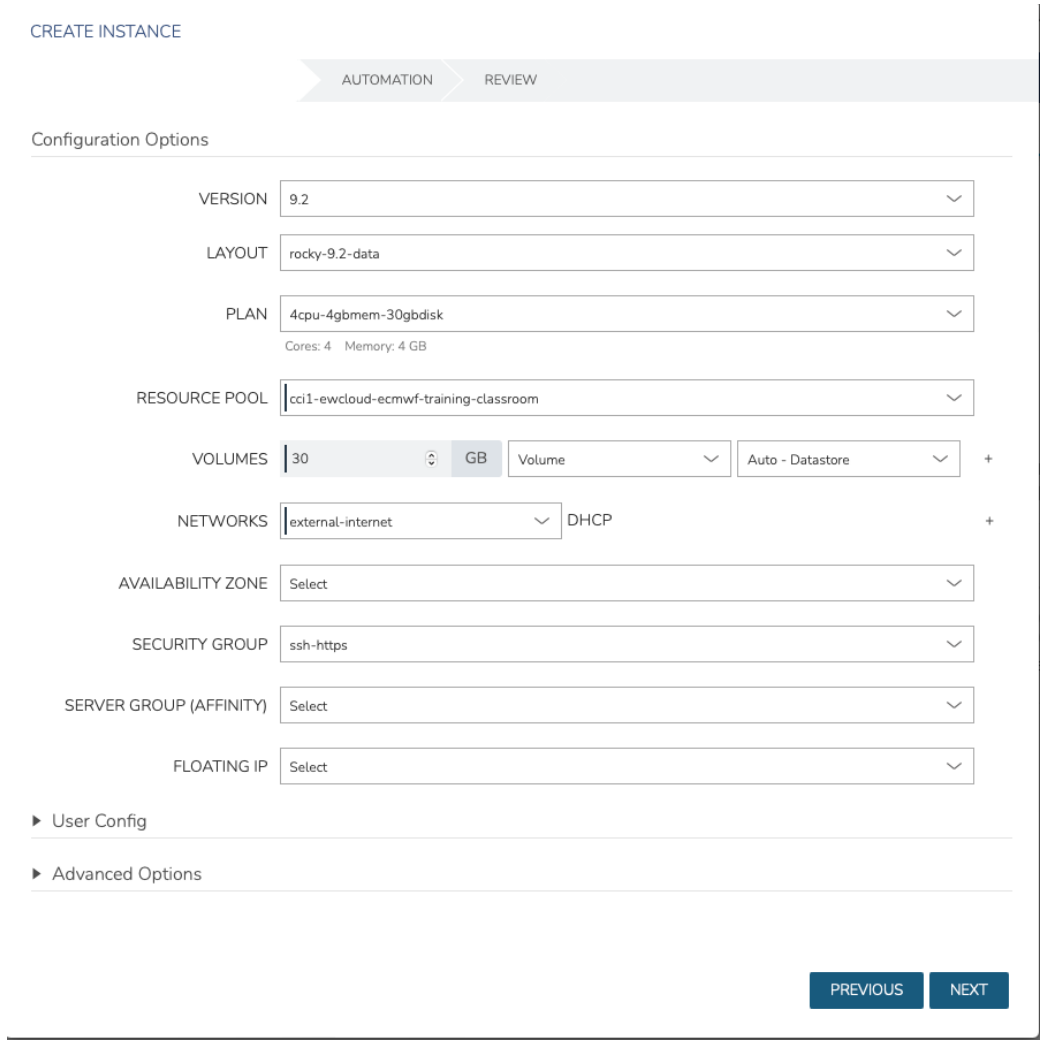

Instance configuration options:

- **Version** - linked to OS version: Rocky Linux 8.x ,9.x and Ubuntu 22.04
- **Layout** – instance configuration flavours :
  - <instance type>-<version>-**data** (pre-installed sw included)
  - <instance type>-<version>-**generic** (plain OS)
  - <instance type>-<version>-**gpu** (GPU drivers included)
- **Plan** – resources configuration : CPU, RAM and local Disk:
  - E.g. : 4cpu-4gbmem-30gbdisk => 4 CPU + 4GB RAM + 30GB disk
- **Volumes** – disks attached to the VM (local and extra disks)
- **Networks**
  - external-internet (public IP)
  - private-<tenancy name> (private network within the tenancy)
- **Security Group** – IP filter rules which define networking access
  - ssh
  - ssh-https
- **Floating IP** (public IP - alternative option to external-internet network!)

Virtual Machines provisioning

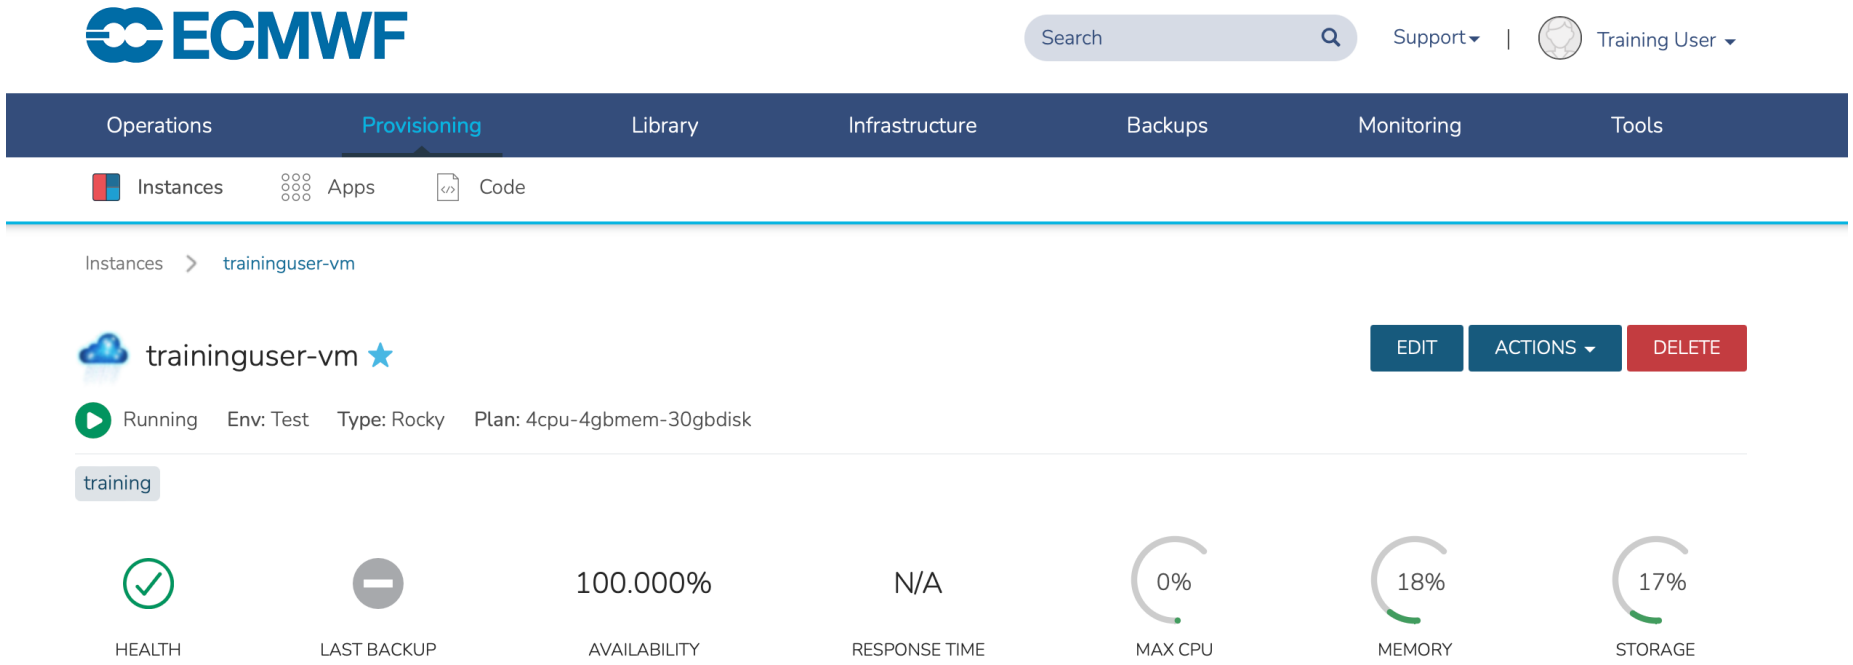

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

S3 Object Storage

- Object storage is a computer data storage architecture designed to handle large amounts of unstructured data
- Data is stored as **objects** within resources called **buckets**
- **Benefits**: high scalability, flat structure, resilience, access protocol
- Supported access via the **S3 RESTful API** is compatible with the basic data access model of the Amazon Simple Storage Service (S3) which runs over HTTPS

https://confluence.ecmwf.int/x/HINyEw

S3 Object Storage - Access

Morpheus GUI

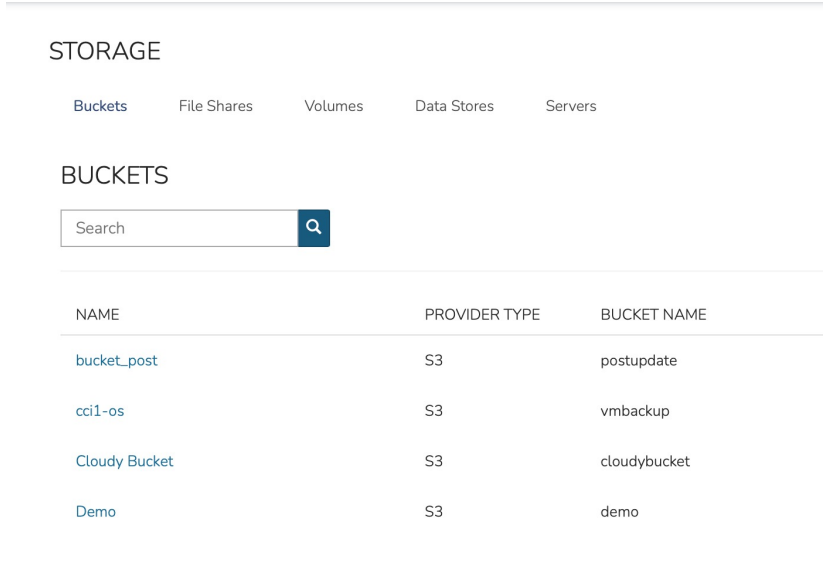

Command Line Tools

- **s3cmd**
- **rclone**
- **awscli**

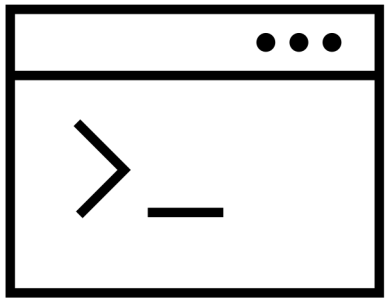

Python Libraries (boto3)

```python
# Initialize the S3 client
s3 = boto3.client(
    's3',
    endpoint_url=S3_ENDPOINT_URL,
    config=Config(
        signature_version=UNSIGNED
    ))
```

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Knowledge Base

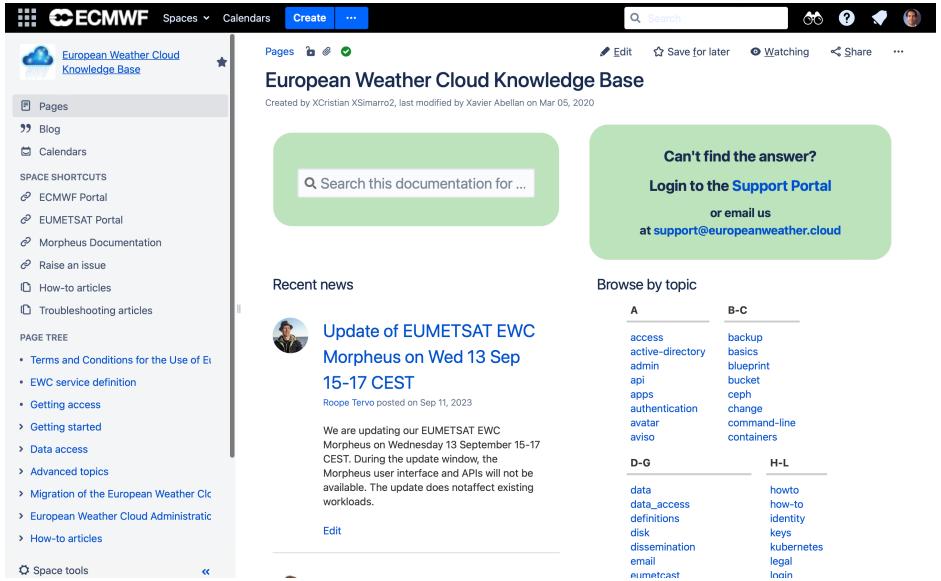

https://confluence.ecmwf.int/x/6J83Cg

KB Content

- Terms and Conditions
- Getting access
- Getting started
- Data access
- Advanced topics
- How-to articles
- Troubleshooting articles
- Discussion platform

Blog Posts

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Support Portal

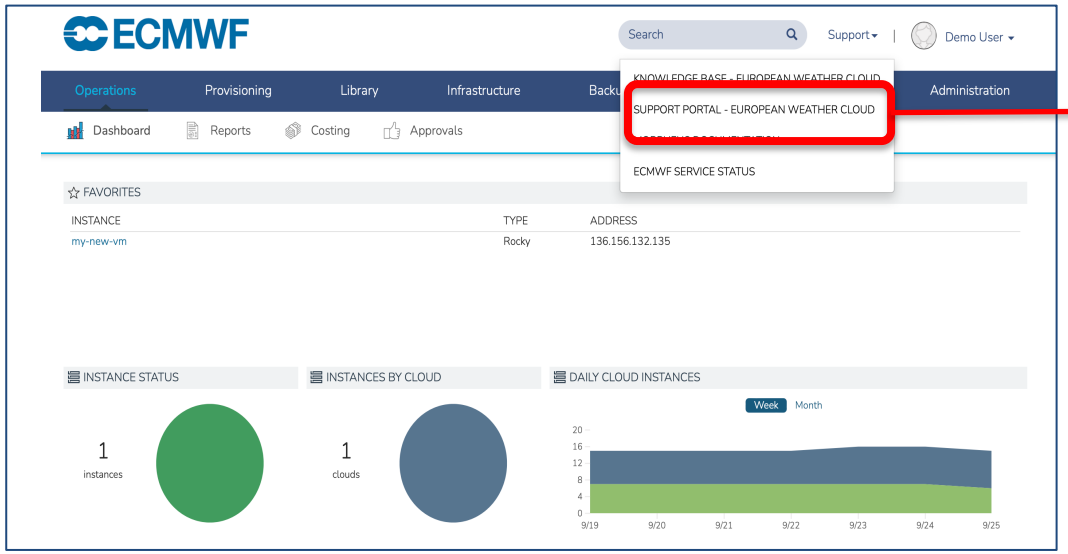

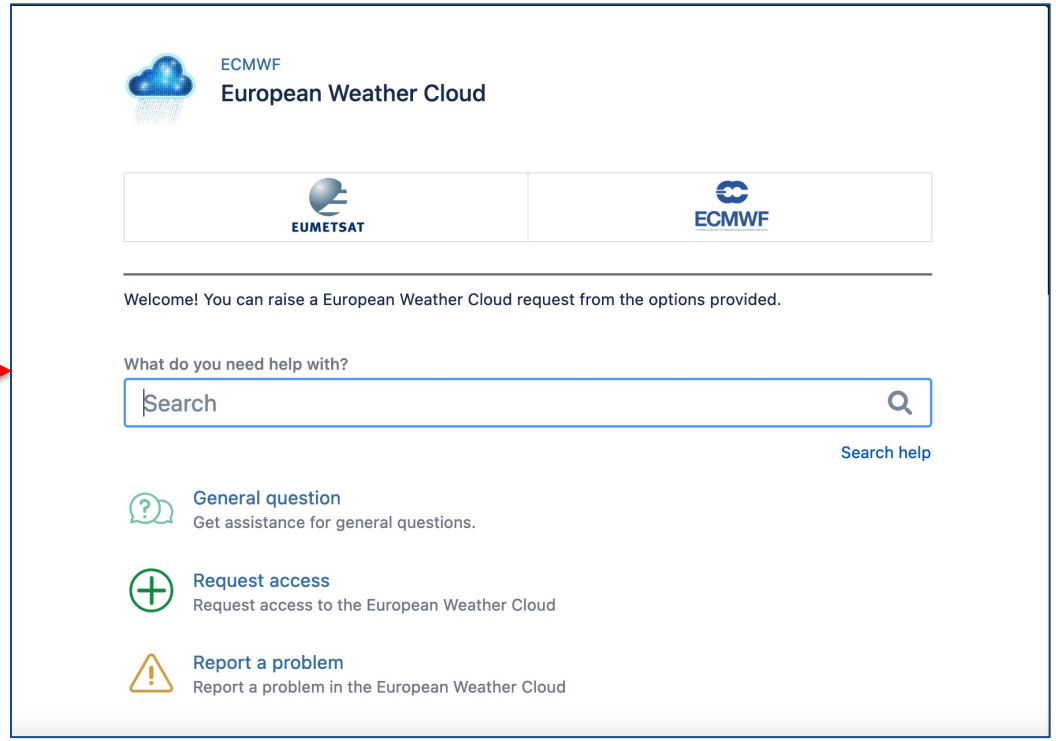

https://jira.ecmwf.int/servicedesk/customer/portal/9

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

EWC Accounting

The **Accounting service** provides a cross-cloud overview of resource usage of the tenancies.

A web GUI provides metrics, time series, graphs, and dashboards displaying the accounting information.

Accounted resources include:

- Virtual Machines
  - vCPU
  - Memory
  - Local Disks
- Object Storage (S3)
- vGPU usage

A **cloud billing unit** is adopted as virtual currency for the accounting of the consumed cloud resources.

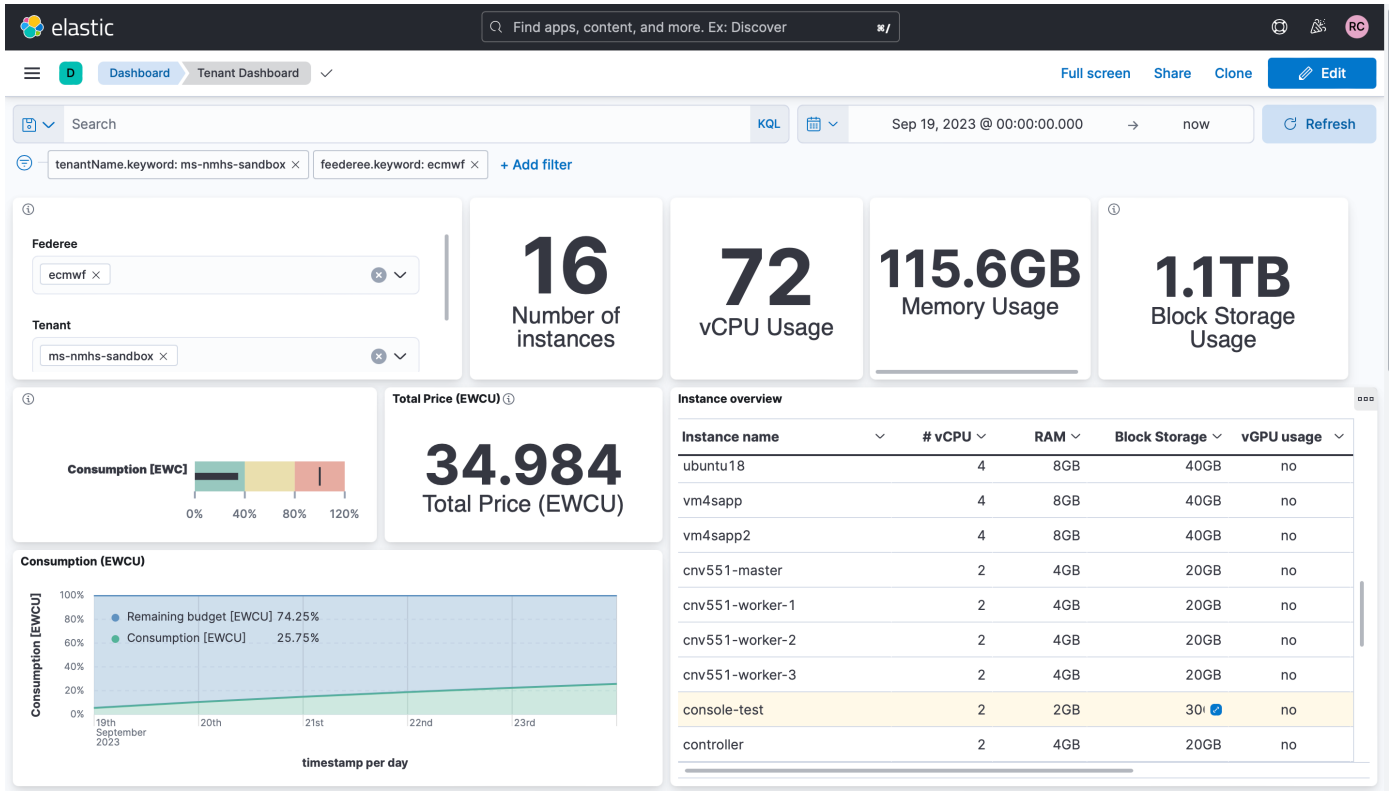

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

EWC Discussion Platform: Rocketchat

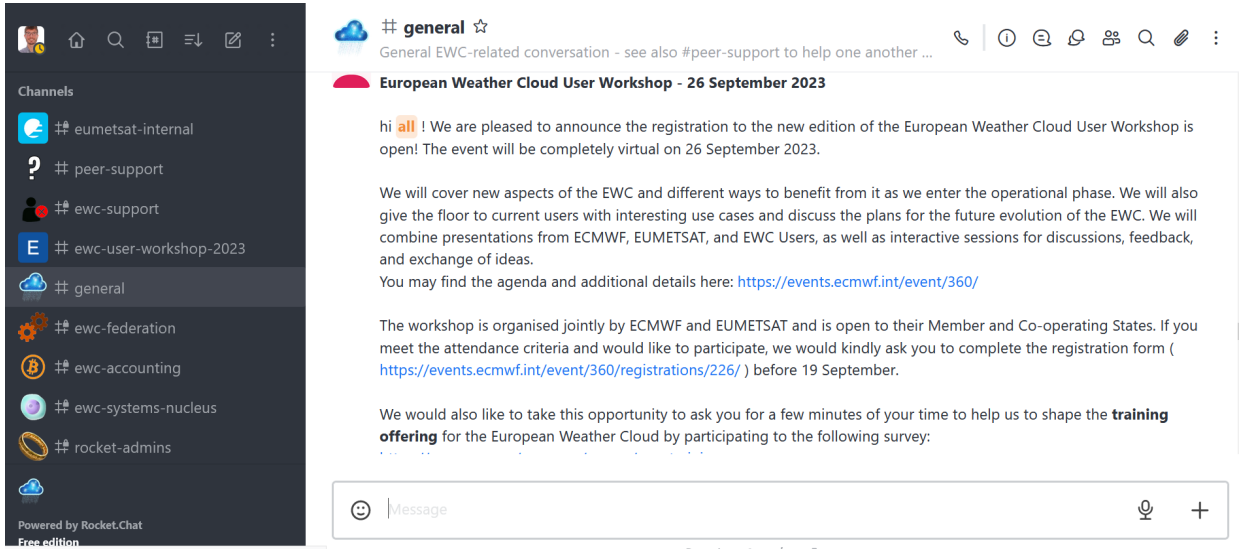

Installation

- Web based
- Desktop App
- Mobile

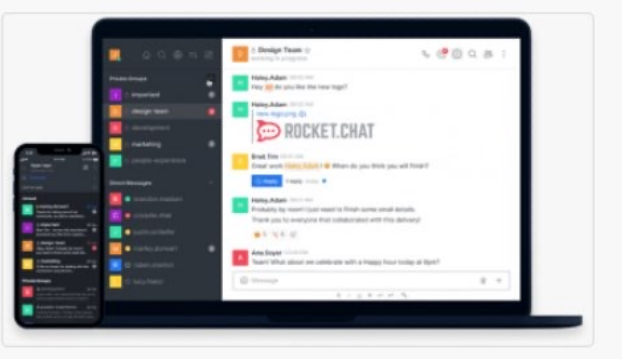

How do I join?

Member and Cooperating States users can self-register (based on email domain): https://chat.europeanweather.cloud

External users can't self-register but can be invited. They need to be accepted by admins.

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Usage context examples

- Pre-operational / Operational usage
- Backup infrastructure
- Cloud “bursting” / elasticity
- Application development and testing support
- Systems architecture setup / testing
- Research / scientific activities
- Training activities
- Collaboration environments

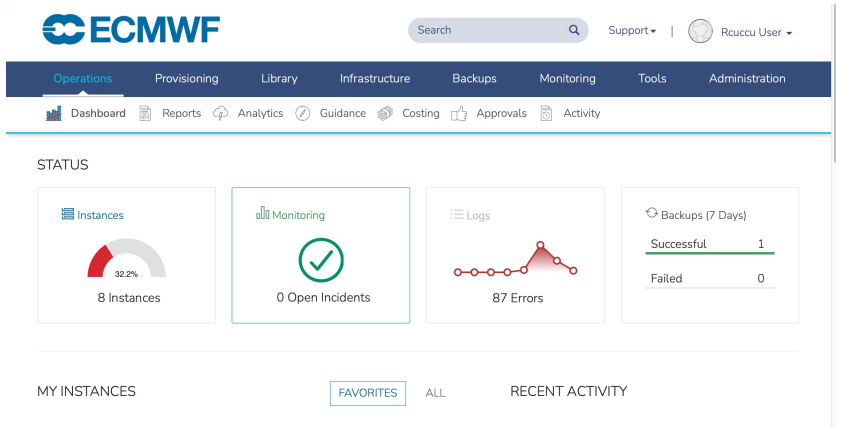

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Training

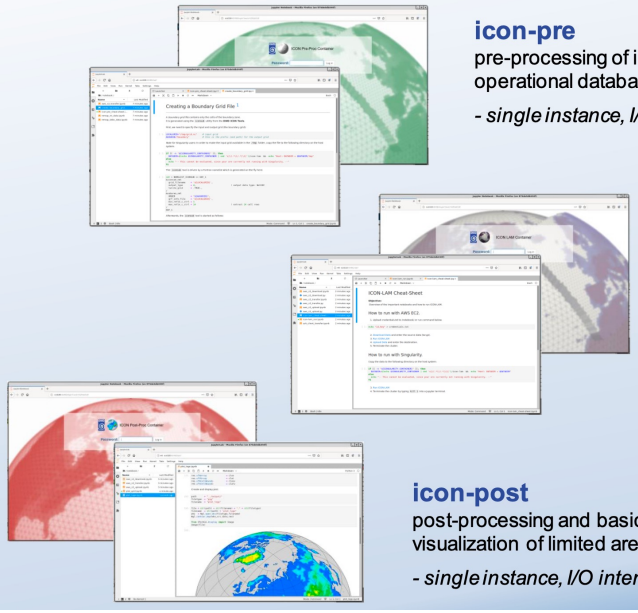

icon-pre
pre-processing of input data from DWD's operational database and other sources.

- single instance, I/O intensive

icon-lam
self-contained, MPI-parallel executable of the ICON limited area model. Ready-to-use for small and medium size setups.

- *"virtual cluster" of multiple instances, CPU+network intensive*

icon-post
post-processing and basic visualization of limited area ICON runs.

- single instance, I/O intensive

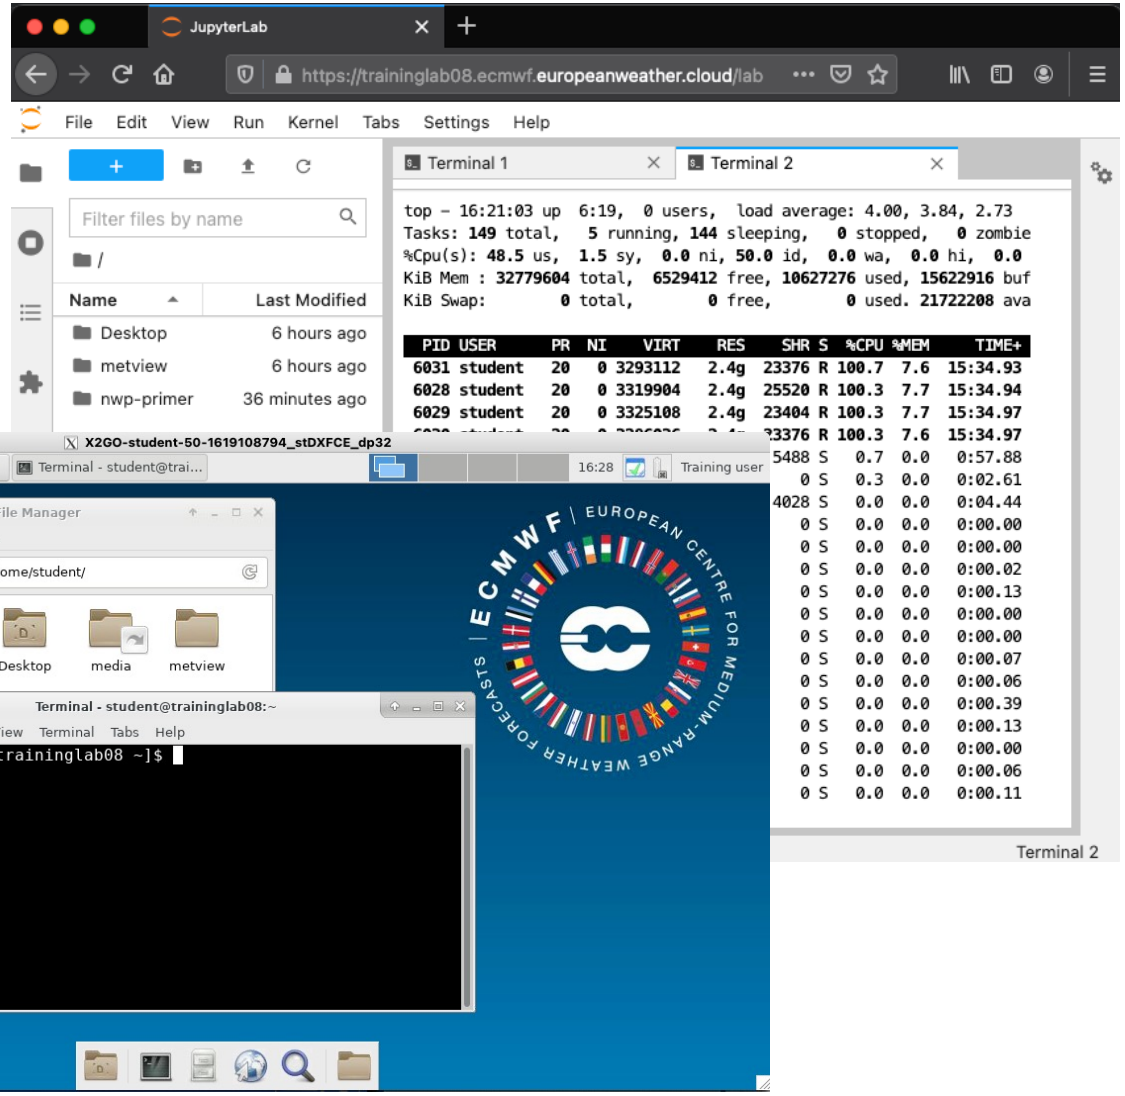

- ICON Lab by DWD
- ECMWF NWP Training labs
- EUMETSAT Training infrastructure

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Artificial Intelligence / Machine Learning

```
In [1]: from keras.optimizers import adam
        import numpy as np
        import nogwd
        import oifsprocess as oip
        from nogwd import gendata,nogwdmodel,dataversion,pressure_weighted_mse

Using TensorFlow backend.

In [2]: (a,b,c,d) = nogwd.loadalldata()

Data version : REDUCED

In [5]: best=nogwdmodel(inputs=a.shape[1],outputs=b.shape[1],
                        verbose=2,hidden_neurons=2048,dropout_alpha=0.0,num_epochs=50,
                        hidden_layers=4,learning_rate=10**(-3.5),batch_size=256,
                        activation='leaky',final_activation='linear',
                        optimiser = adam,loss=pressure_weighted_mse,model_config=3027
                        )

In [6]: best.load_model()

3027

In [7]: best.model.evaluate(c,d)

257184/257184 [==============================] - 17s 67us/step

Out[7]: 0.023816158334751605

In [12]: import nogwd_plots
         import imp
         imp.reload(nogwd_plots)
         nogwd_plots.timepredictionquality(best,np.arange(1500,30001,250))

<Figure size 432x288 with 0 Axes>
```

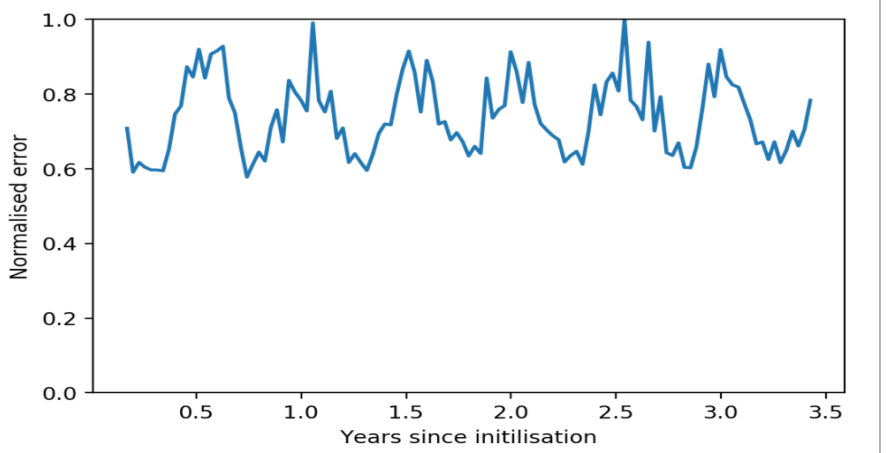

- Oxford University & ECMWF collaboration
- ECMWF S2S AI/ML competition
- ESoWC/Code4Earth projects
- Exploratory activities by AEMET and DWD
- DWD Nowcasting with machine learning
- Italy post-processing with machine learning
- FUSEDCAST nowcasting using MTG/GOES data

TensorFlow

Keras

PyTorch

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Data proximate computing in EWC with Dask

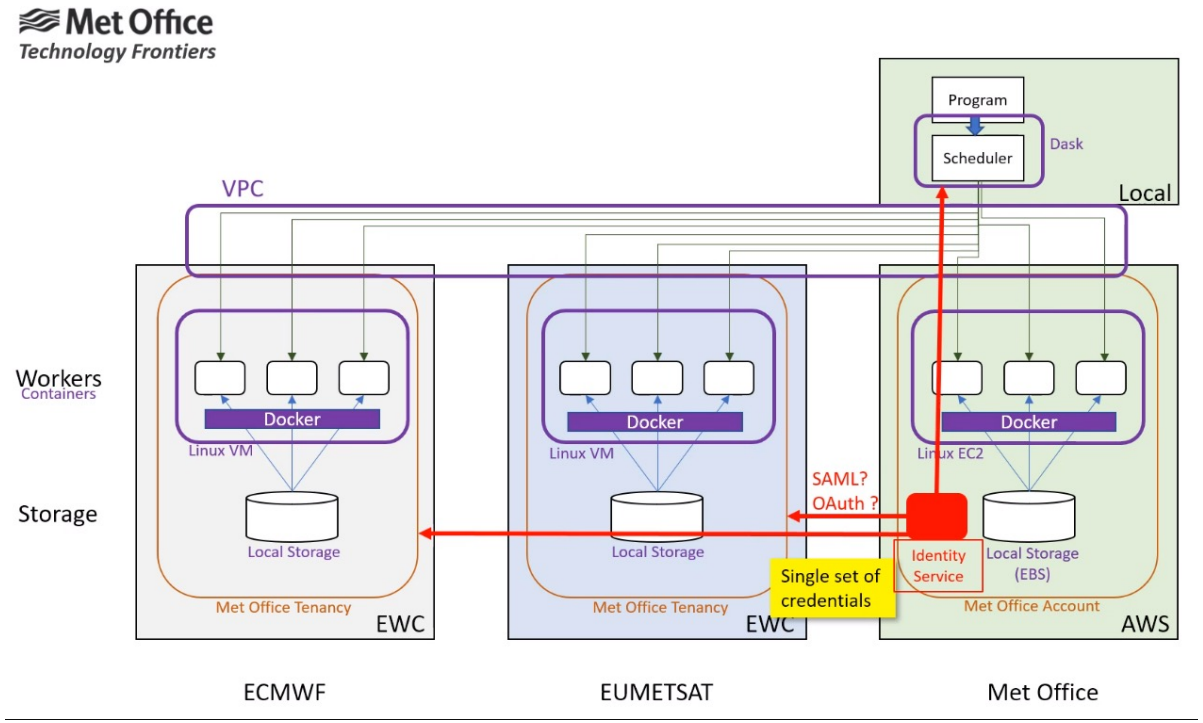

Image: Met Office

```
with in_org('metoffice'):
  dataset =
xarray.open_dataset('metoffice.nc').c
hunk({"latitude": 10})

with in_org('eumetsat'):
  comparison_dataset =
xarray.open_dataset('eumetsat.nc').ch
unk({"latitude": 10})

averages =
dataset.mean('realization',
keep_attrs=True)
diff = averages.isel(height=5) -
comparison_dataset
diffed =
diff.compute(optimize_graph=False)
```

Example simplified from: https://dask.discourse.group/t/understanding-work-stealing/335/9

For more information:

https://blog.dask.org/2022/07/19/dask-multi-cloud

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Data & visualisation services

- Geoportal Service by DWD
- Climate Explorer by KNMI
- ESSL Displayer
- SmartMet Server, data server by FMI

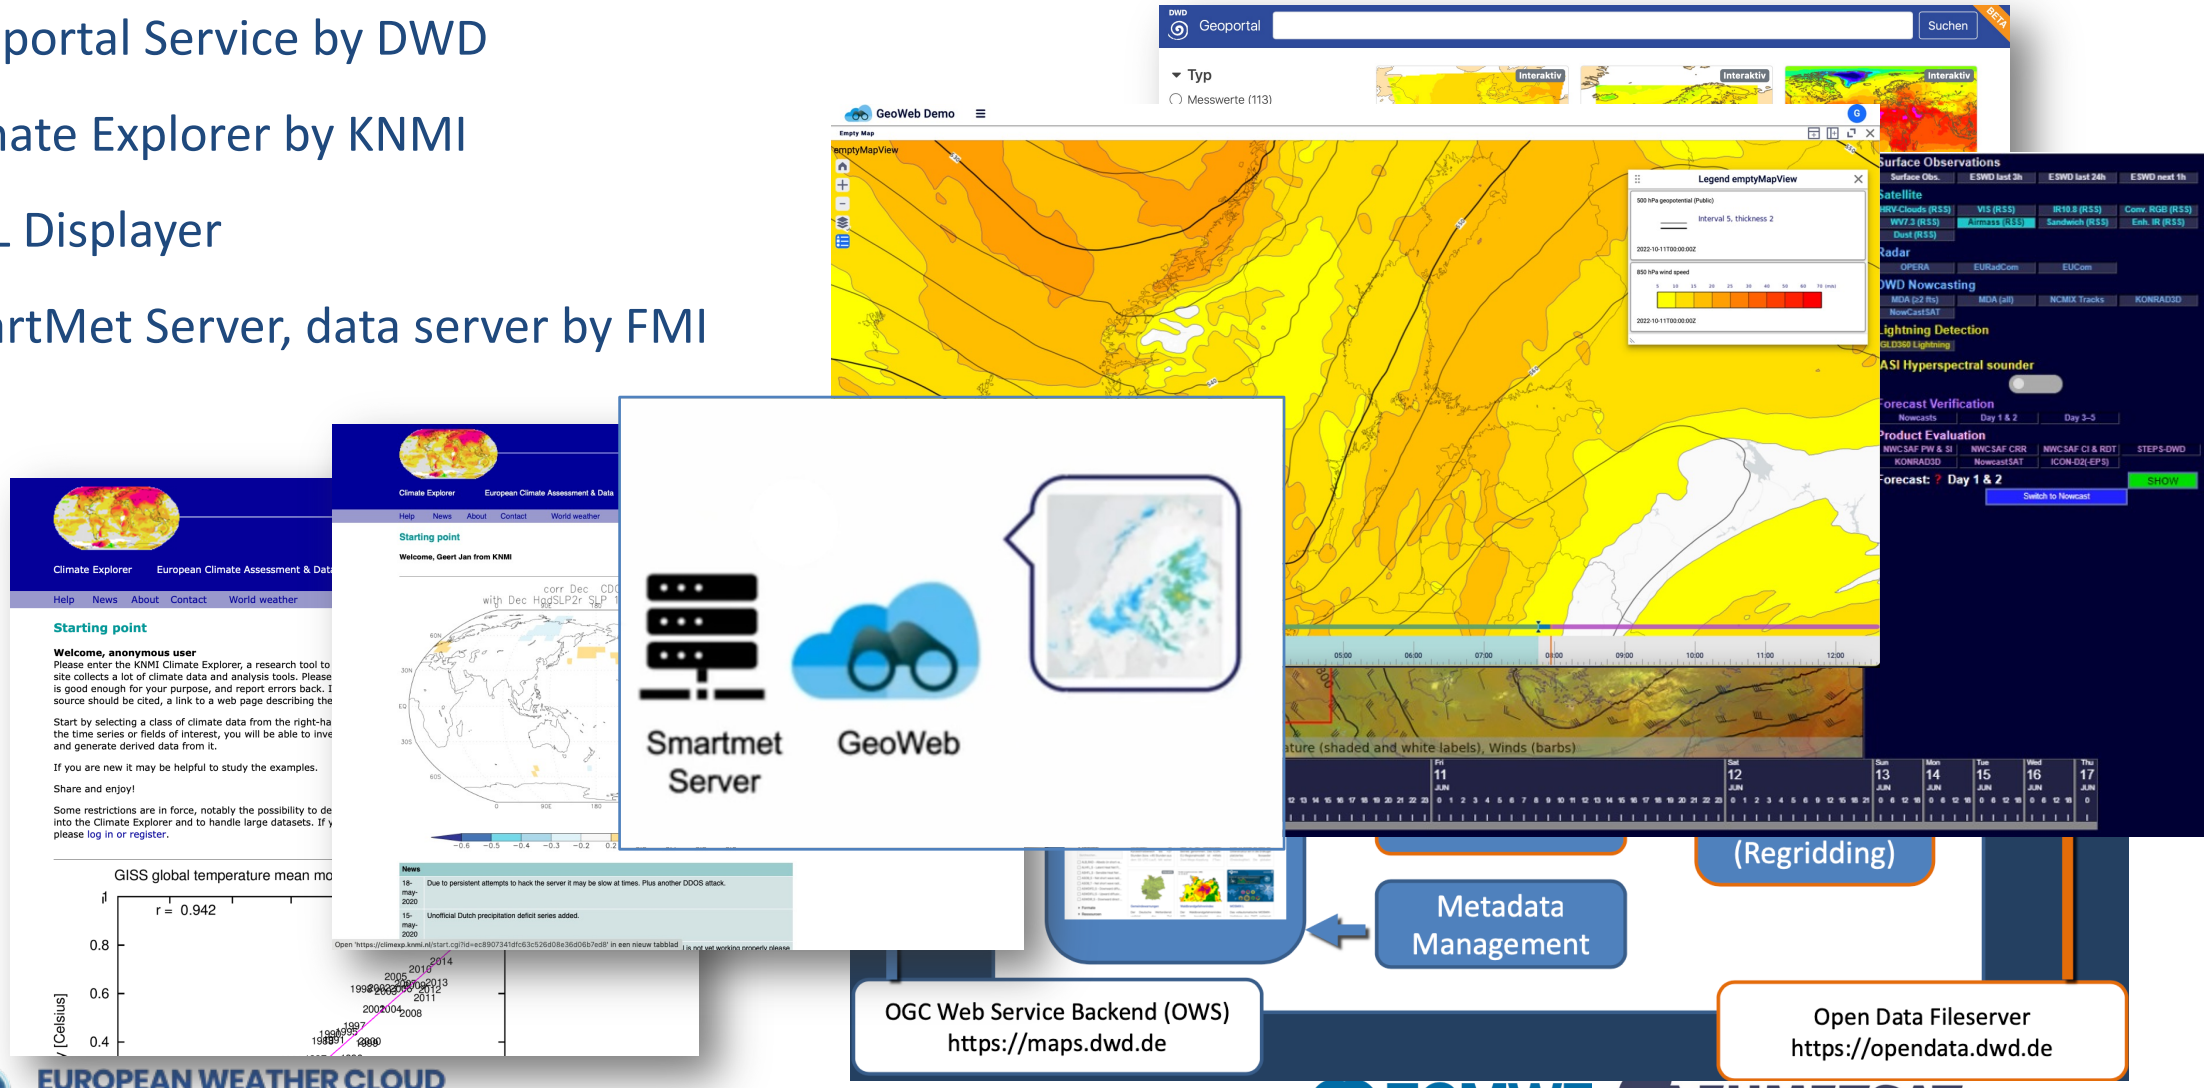

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

International collaboration projects

- SEE-MHEWS-A
- NordSat PyTroll MTG product generation

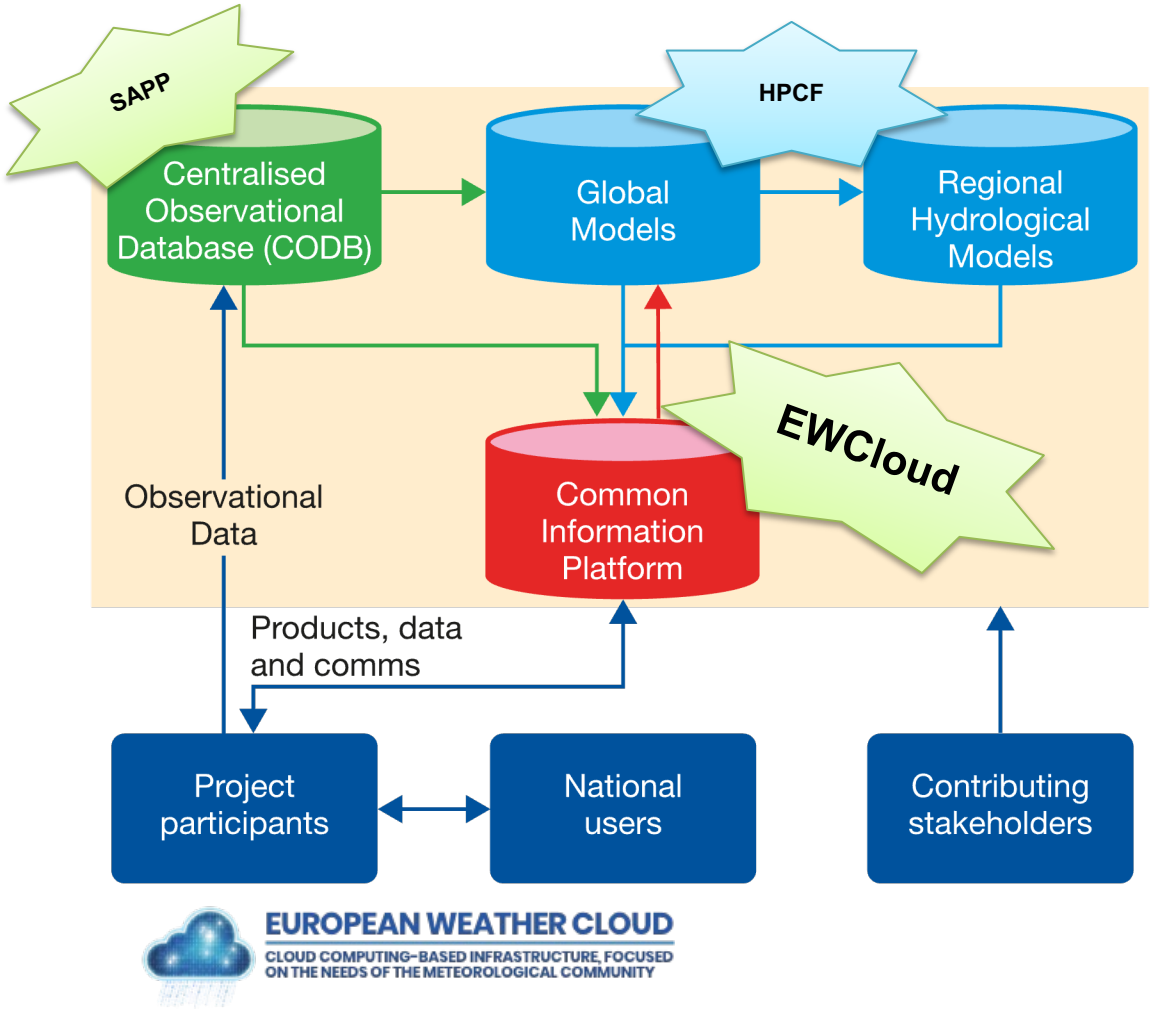

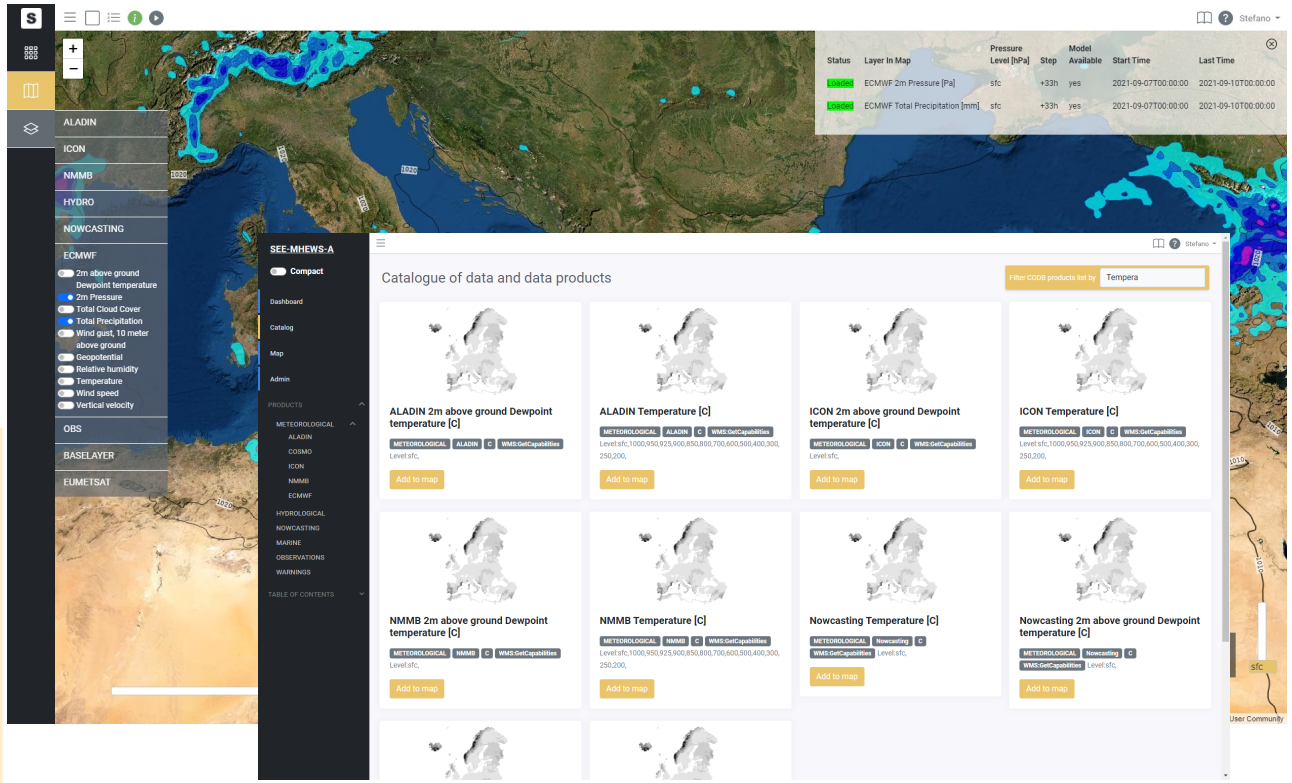

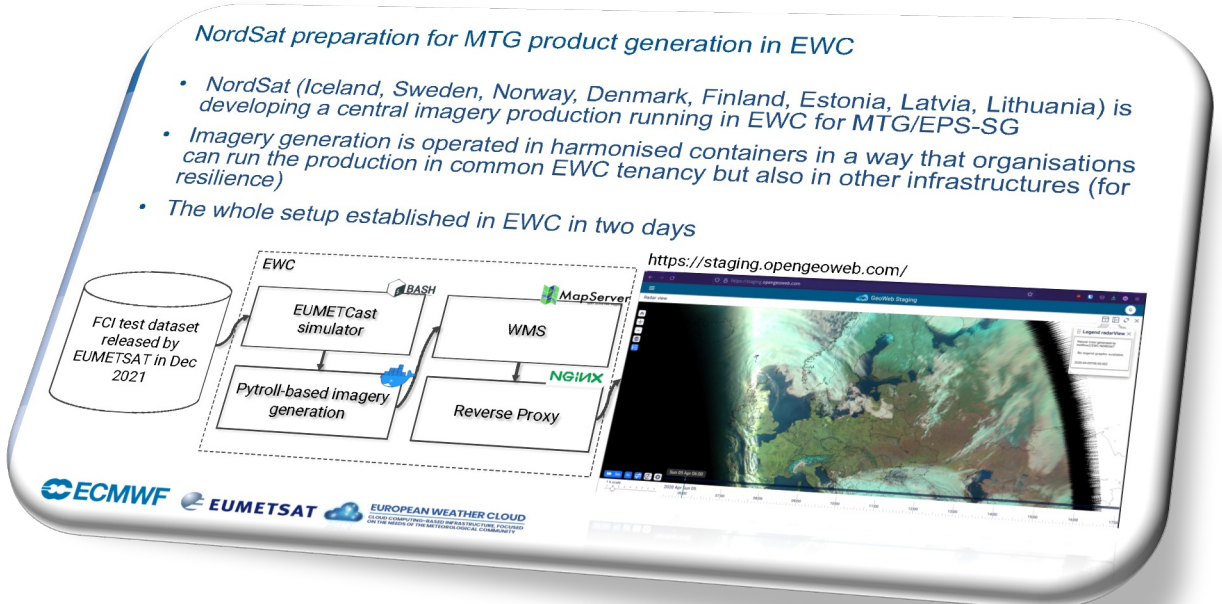

Contact Us and Request Access

www.europeanweather.cloud

Contact Us

support(at)europeanweather.cloud

EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

ECMWF EUMETSAT

Questions?