Status and summary of the Pilot

European Weather Cloud User Workshop

November 10-11, 2021

Roberto Cuccu (ECMWF) , Roope Tervo (EUMETSAT)

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

European Weather Cloud

www.europeanweather.cloud

- Three-year pilot project started in January 2019 -ECMWF and EUMETSAT
  - Basic goal is **to bring the computation resources (Cloud) closer to the data** (meteorological archive and satellite data)

- The Pilot project includes:
  - Building infrastructure
  - Hosting and supporting use cases
  - Gather feedback and prepare for next phases

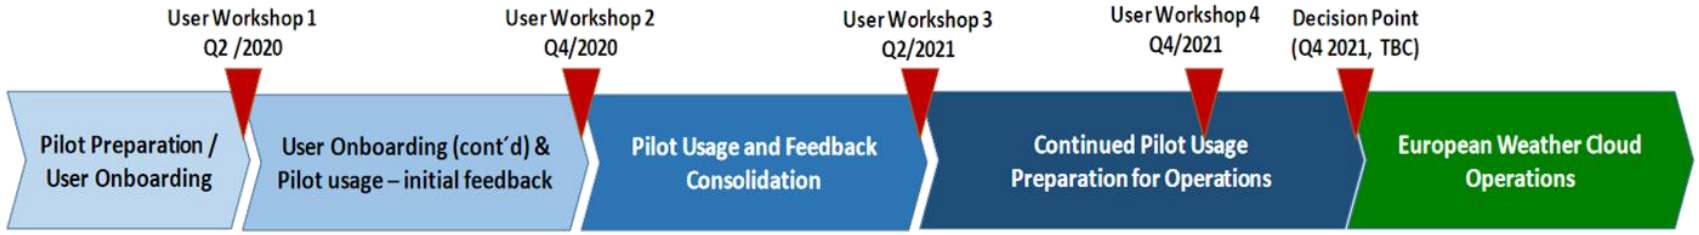

ECMWF EUMETSAT EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

High Level Design and users entry point

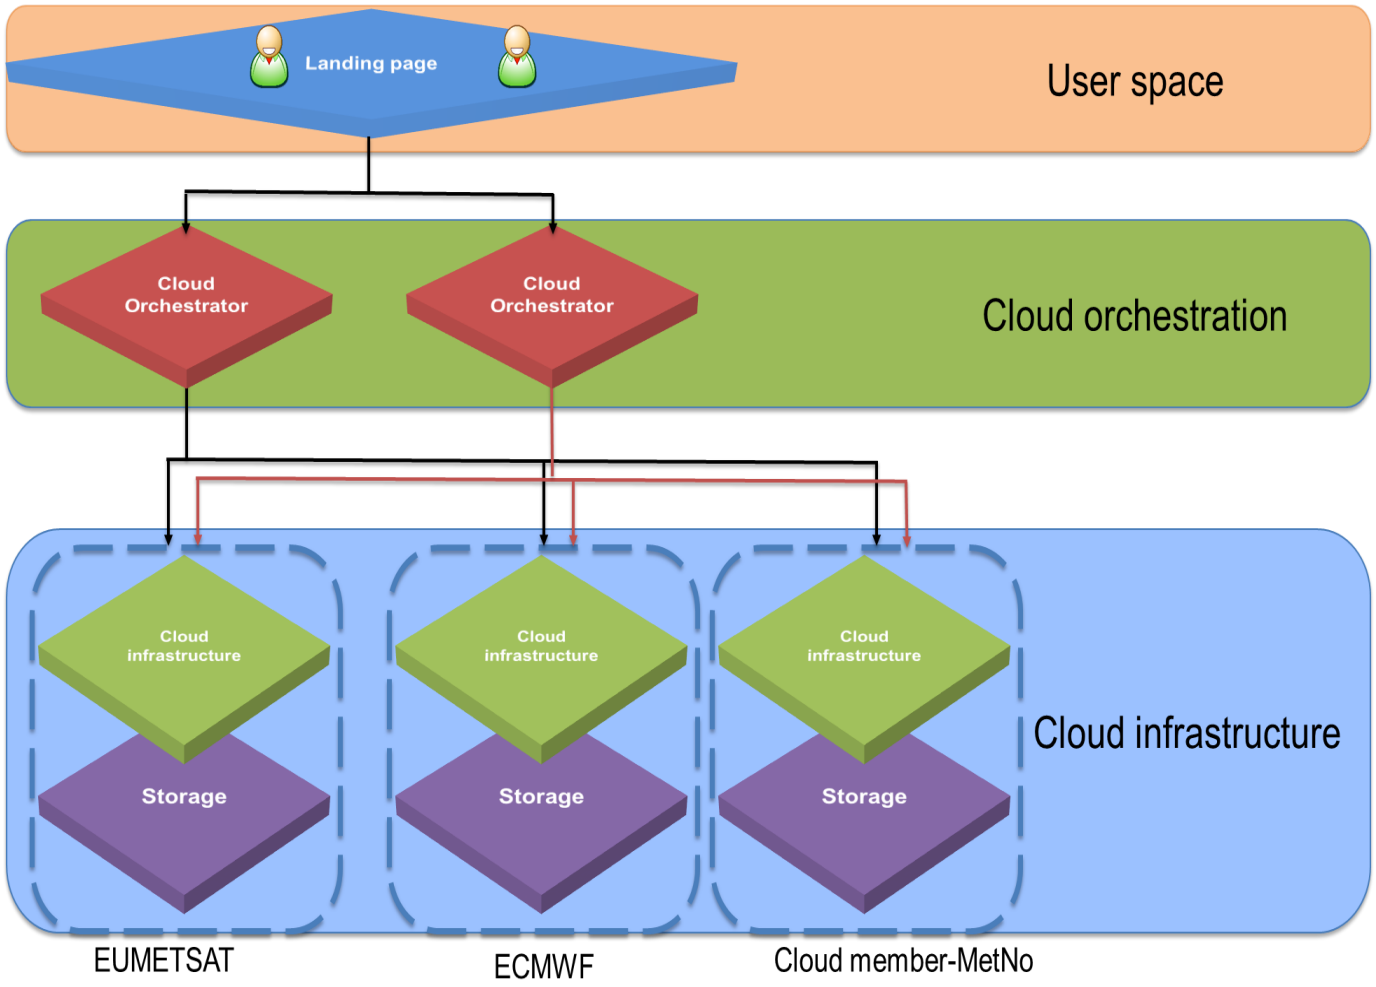

MORPHEUS

openstack®

Underline cloud infrastructures not directly accessible by the users

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

Cloud Management Orchestrator - Morpheus

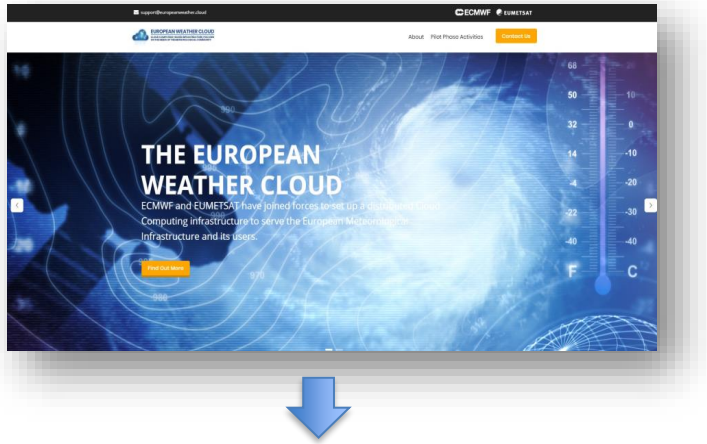

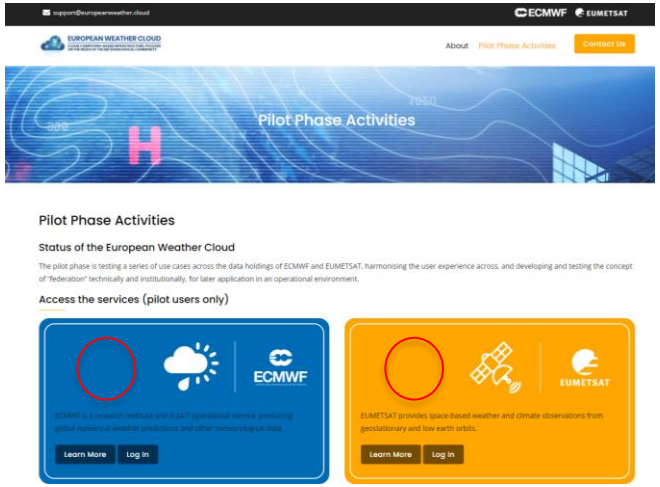

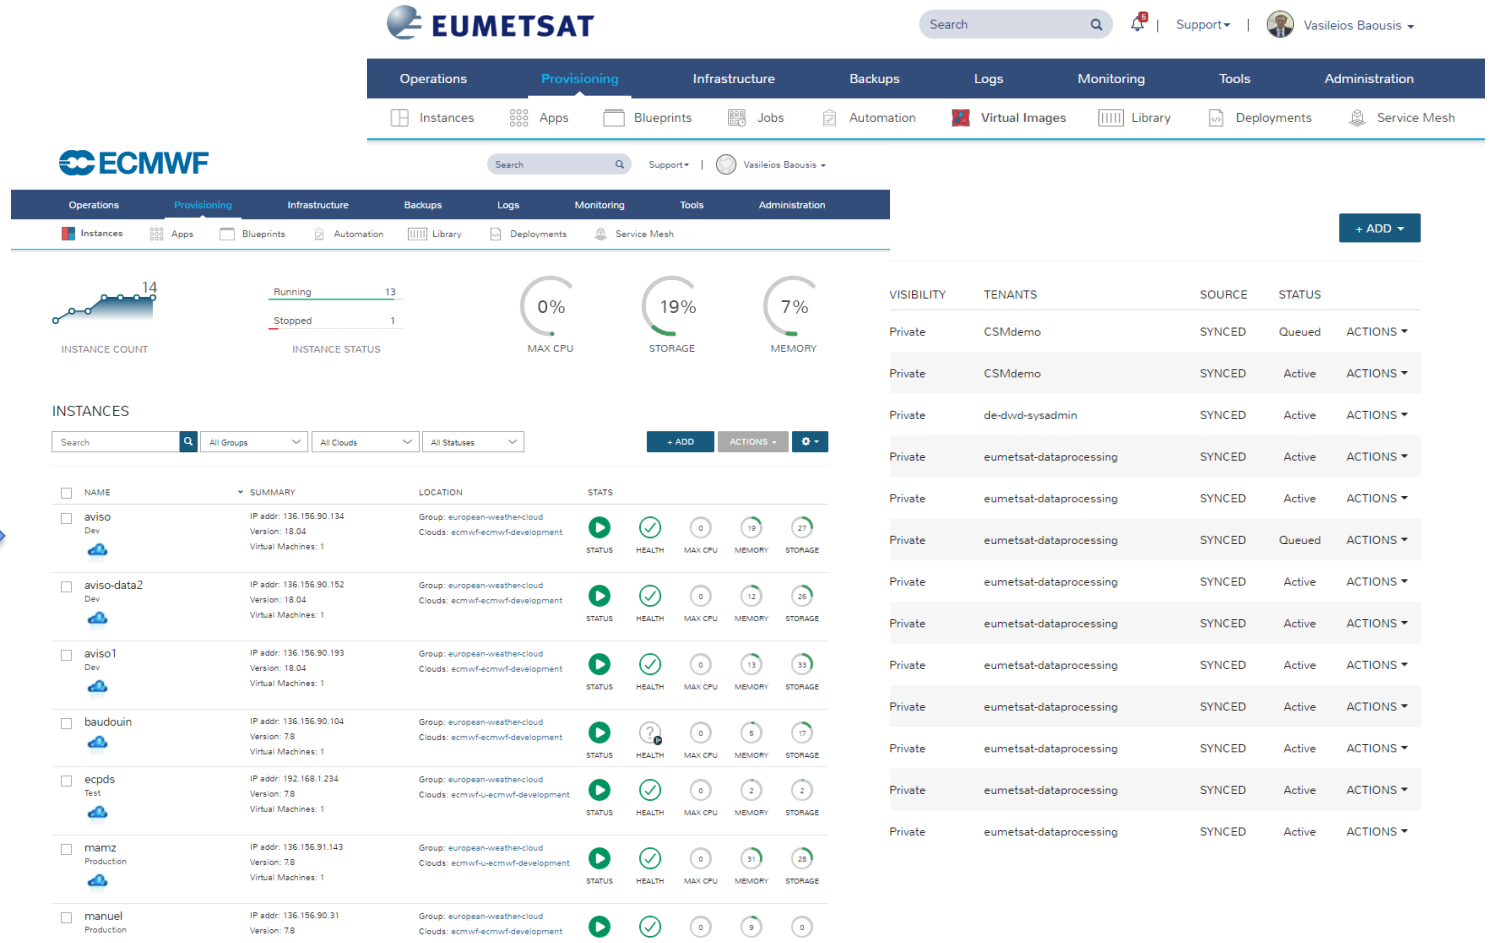

European Weather Cloud Pilot infrastructure

ECMWF

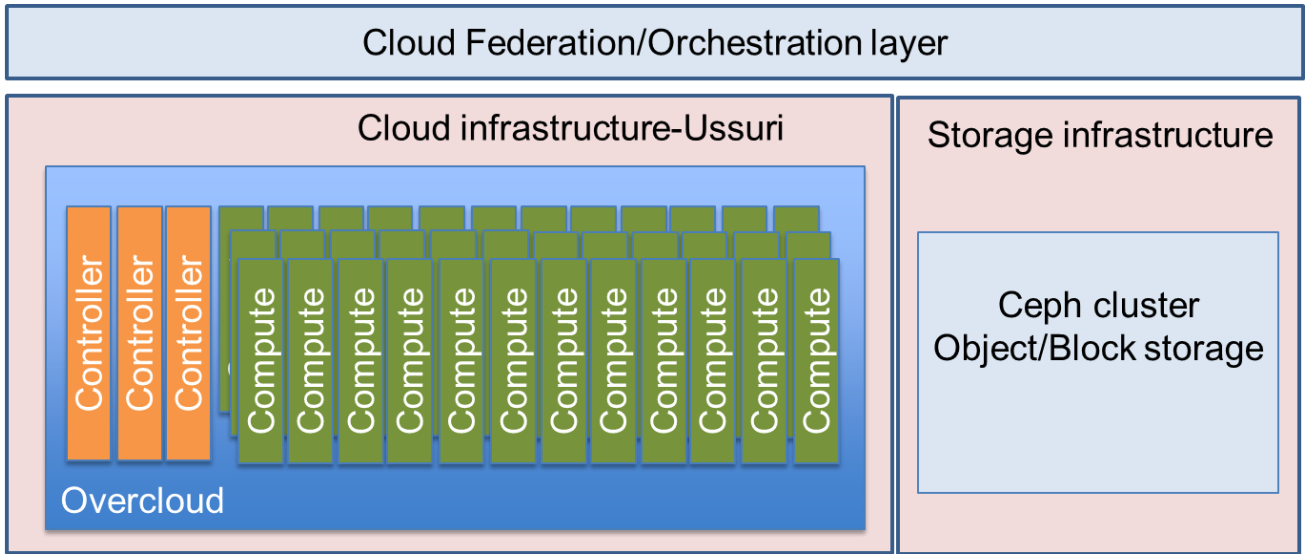

EUMETSAT

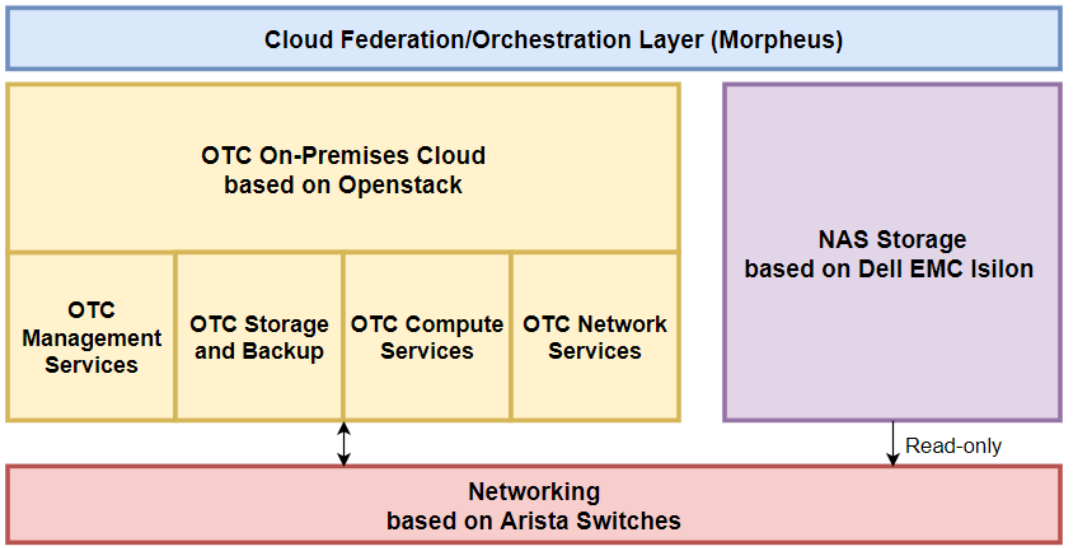

Support services

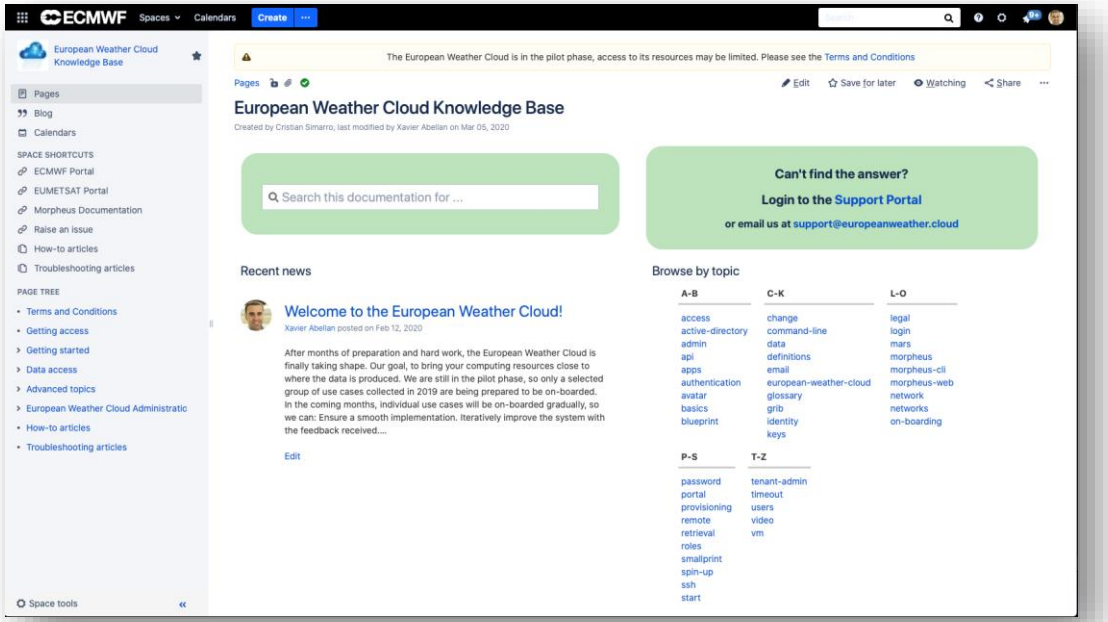

Knowledge Base

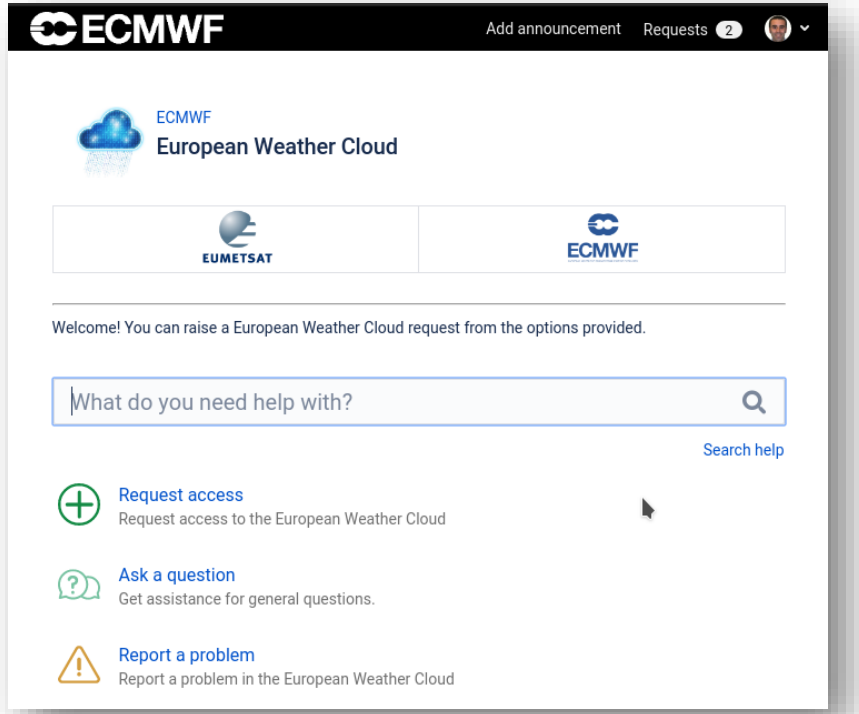

Support Portal

ECMWF EUMETSAT EUROPEAN WEATHER CLOUD
CLOUD COMPUTING–BASED INFRASTRUCTURE, FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

Cloud offer (currently IaaS model in the Pilot)

Cross-clouds resource provisioning

Self provisioning of customisable Virtual Machines (CentOS and Ubuntu)

VM templates with pre-installed software

”Fast-track” data access

Self provisioning of block storage volumes and S3 object storage usage

GPUs instances

Self provisioning of Kubernetes clusters

Cloud orchestrator APIs usage

Integrated automation (e.g. via Ansible)

…and more

44 different kind of use cases in the pilot phase (which of 44 were useful)

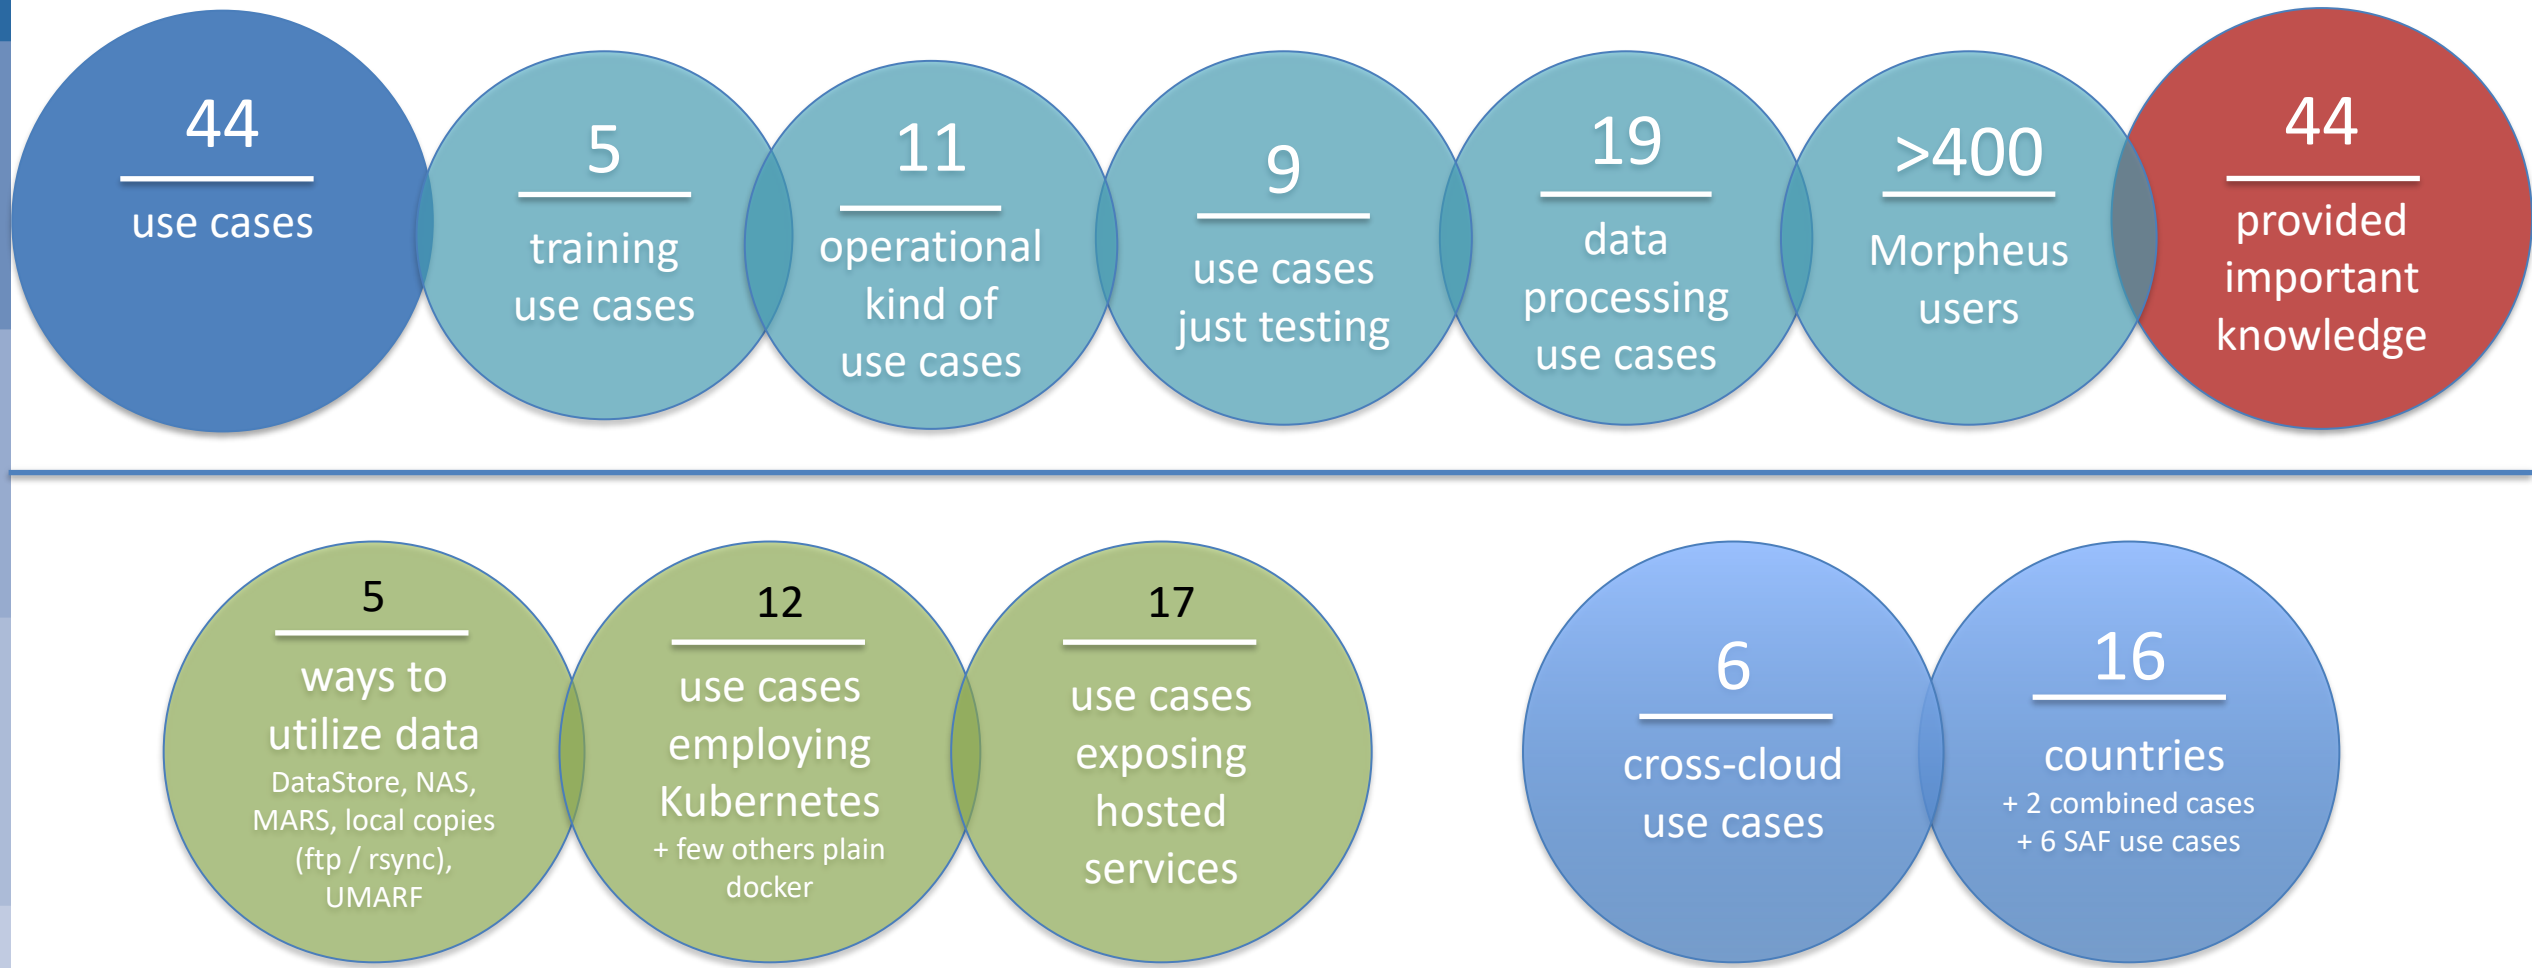

Users Feedback

Feedback collected via: users workshops, surveys, support tickets, monitored cloud usage

Some major points reported so far:

- Integration with other facilities (e.g. HPC)
- Containers and kubernetes clusters
- Higher level services and predefined building blocks
- Data access and availability
- IT Security and data protection

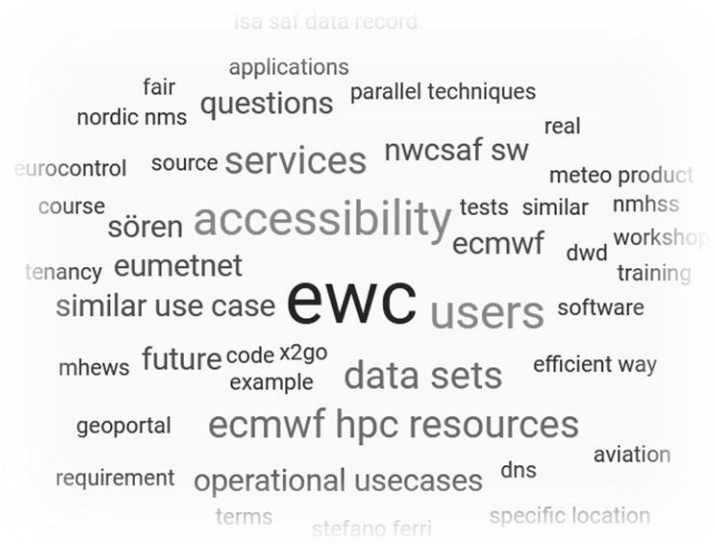

EWC Pilot - Survey Results

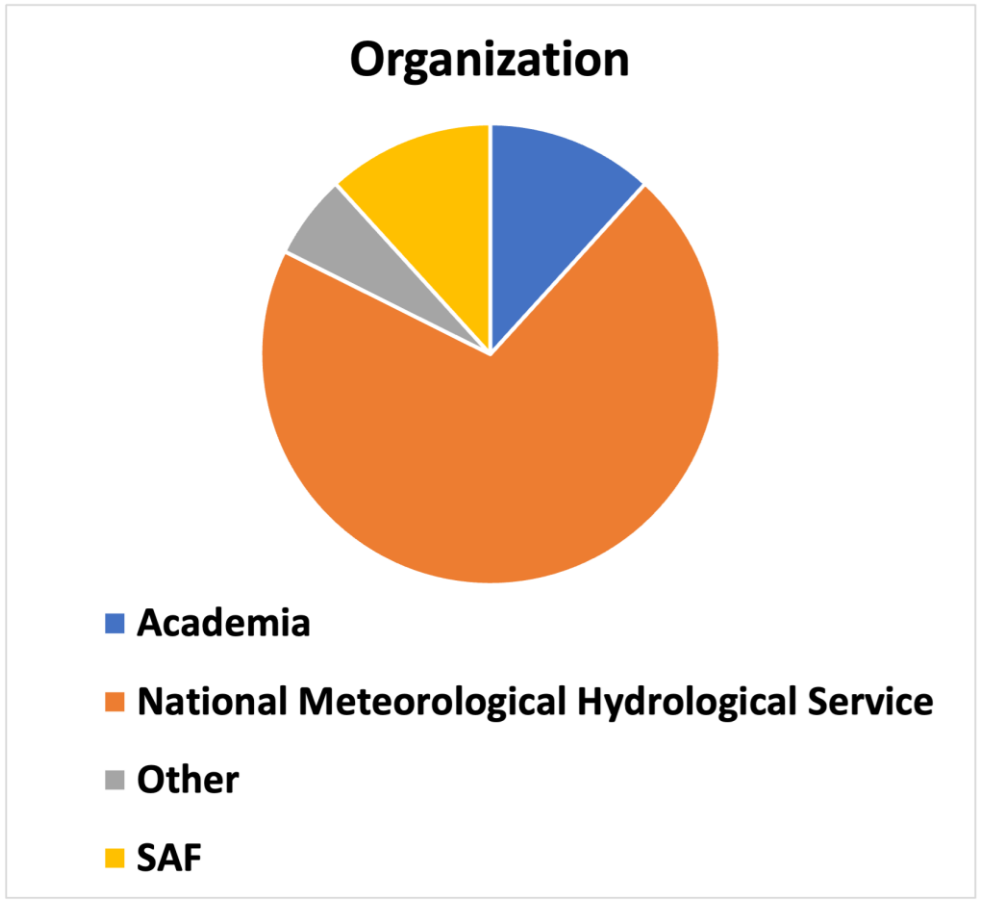

Use or intended use of European Weather Cloud?

Operational systems	
Batch processing (reprocessing, post-processing/extracting subsets of a dataset, etc..)	
Providing a remotely accessible service (e.g. website / web tools, online visualisation)	
Interactive work / collaborative environment	
Training environment	
Machine learning	
Experimentation and exploration of cloud technologies	

Survey Results – Pilot phase

Positive feedback:

- Getting support (initially or for technical help), via Knowledge Base, Support Portal, etc.
- Getting a tenancy
- Exporting results and data out of the cloud

Need improvements:

- Understanding how to use cloud resources for your use case
- Setting up virtual machines, clusters, or any other virtual infrastructure
- Accessing data

Obstacles:

- Insufficient resources (e.g. wanted GPUs, RAM or disk quota too small)
- Technical knowledge and system administration

Survey Results – Next Phases

- Confirmed interest in accessing cloud resources after Pilot phase
- High importance in having access to training courses and workshops on the use of the cloud

Preferred ways to access to the operational European Weather Cloud:

Ask for a share of the national contingent (via Member State Computing Representative)	
Apply for an upcoming EUMETSAT R&D call / ECMWF Special Projects	
Not sure, I'd like to discuss options	

We are looking forward the operational phase

- Carry-ons from the pilot phase
- New and interesting uses we haven't seen yet
- Operational applications
- Projects with longer temporal scope