ECMWF Feature article

from Newsletter Number 180 – Summer 2024

COMPUTING

The European Weather Cloud is now operational

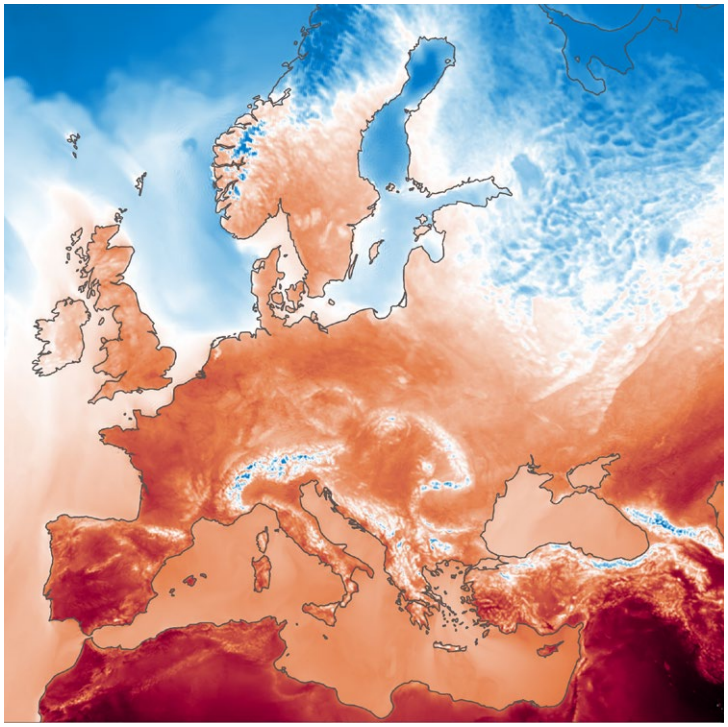

www.ecmwf.int/en/about/media-centre/media-resources

doi: 10.21957/xawe573ul9

This article appeared in the Computing section of ECMWF Newsletter No. 180 – Summer 2024, pp. 37–43.

The European Weather Cloud is now operational

Xavier Abellan, Vasileios Baousis, Ricardo Correa, Roberto Cuccu, Cristina Duma, Charalampos Kominos, Samuel Langlois, Umberto Modigliani

The European Weather Cloud (EWC) is a community cloud computing platform jointly operated by ECMWF and the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT). The concept was proposed by ECMWF, building on its extensive experience in providing computing access to its Member and Co-operating States, for example through the ecgate service. The EWC facilitates data access, increases data processing capability, and fosters new forms of collaboration among national meteorological and hydrological services and researchers in ECMWF and EUMETSAT Member and Co-operating States. The initiative started in 2019 with a pilot phase laying the foundation of the current service. The pilot phase had several objectives, including: i) establishing the overall governance and coordination processes among the federating partners; ii) setting up and testing a cloud infrastructure and gaining important operational experience; and iii) on-boarding and supporting a variety of use cases covering different thematic areas and applications in order to gather feedback and requirements to continue to develop and evolve the platform. In 2023, the operational ECMWF cloud infrastructure was deployed in the data centre in Bologna (Italy), co-located with the other ECMWF systems. Thereafter, the migration of existing users, from the pilot cloud infrastructure running in the original Reading (UK) data centre to the operational system in Bologna, was performed. The EWC was declared operational on 26 September 2023.

What is the EWC?

According to the original vision, the EWC is "the cloud-computing-based collaboration platform for meteorological applications development and operations in Europe enabling the digital transformation of the European Meteorological Infrastructure".

The EWC brings together users from different countries and organisations to collaborate and share resources. It allows users to customise and deploy their applications and workflows and to build and expose services on the Internet. This flexibility in the deployment of user applications and workflows is one of the main benefits of the EWC compared to a managed high-performance computing facility (HPCF) service. Users run applications and services next to where the data is produced, avoiding large data movements over the network.

Who can use the EWC?

The EWC service is mainly available to eligible users from the European meteorological community. In particular, it is accessible by:

- users from the national meteorological and hydrological services (NMHSs) and public institutions in ECMWF and EUMETSAT Member and Co-operating States
- research users in the context of ECMWF Special Projects and EUMETSAT R&D calls
- ECMWF and EUMETSAT for internal use in the scope of their mandates
- members of the European Meteorological Infrastructure (EMI), including organisations such as EUMETNET
- ECMWF Third Party Activities and Optional Programmes approved by ECMWF's Council
- users from the NMHSs of the World Meteorological Organization (WMO) and research organisations for activities aligned with ECMWF's mission.

Figure 1 illustrates the availability of the EWC.

What are the main features of the service?

The platform provides many benefits to the user community, including access to a cloud computing facility, flexibility in the provisioning and management of the deployed resources and, not least, all the advantages of a community environment, such as shared knowledge, experience, applications and data, as well as synergies and collaborations.

 doi: 10.21957/xawe573ul9

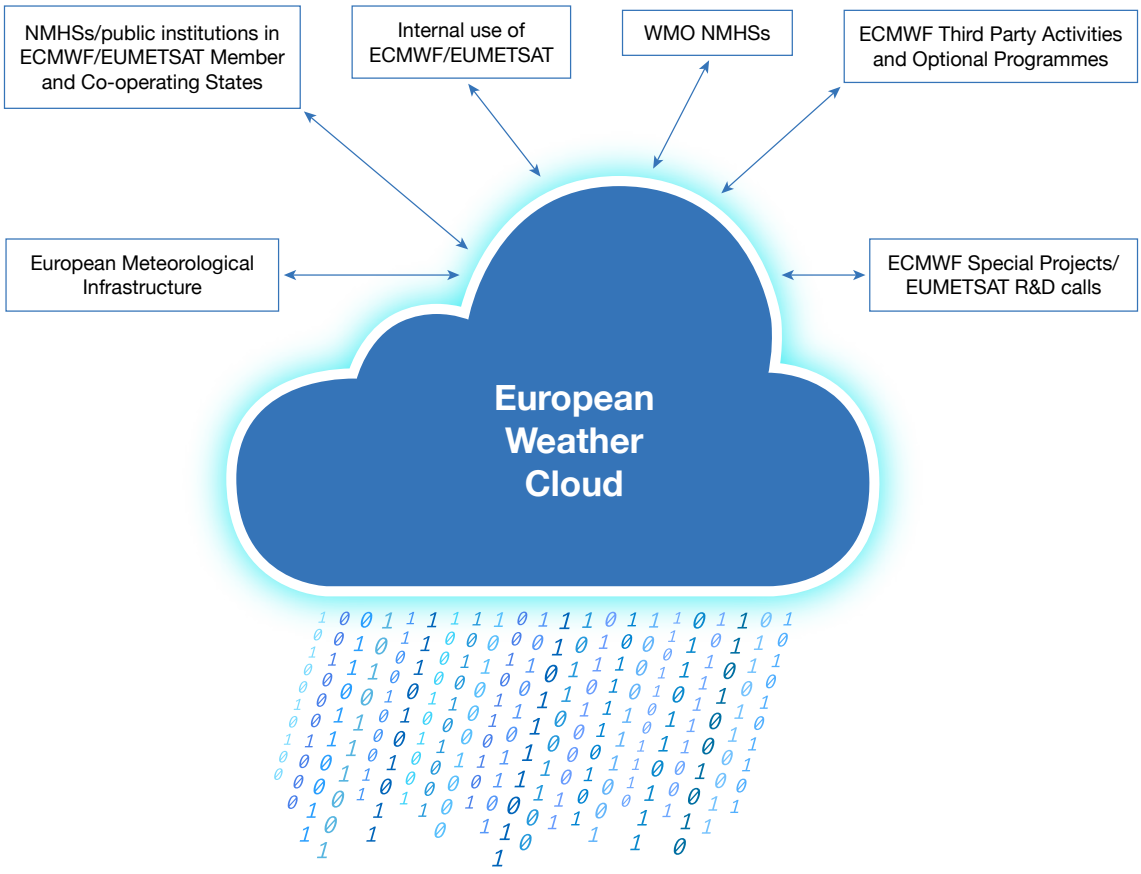

Figure 1 The European Weather Cloud service is available to different groups of users. For more details, see the text.

An important aspect of the platform is optimised access to the data repositories of the two organisations, which include services for the tailored retrieval and dissemination of petabytes of meteorological forecasts as well as climate, ocean, and satellite data.

The EWC provides users with the capability to access cloud computing infrastructure from ECMWF and EUMETSAT by deploying custom virtual environments. These range from virtual machines to more complex setups, including tailored storage and networking. Users can automate the management of resources with the help of blueprints, templates, and other automation tools to ensure the reproducibility of their workflows. Additional services are at the user's disposal for monitoring, reporting, and accounting of the provisioned resources.

Users can manage their assigned virtual environments by using a cloud orchestrator tool that enables the control and deployment of resources, workflows or services on the underlying cloud infrastructures operated at ECMWF and EUMETSAT.

Finally, a set of support and collaboration tools, such as a support portal, shared documentation, and a discussion platform, complements the solutions provided. All these services and features are constantly evolving following an established roadmap, taking into account users' feedback.

How is the EWC integrated with other services at ECMWF?

By running in the same ECMWF data centre, the EWC can leverage the co-location of its infrastructure together with the rest of the service portfolio operated by ECMWF.

In particular, a 'fast-track' data access to ECMWF data services brings the computation and services close to the data sources. Users can exploit a combined set of retrieval and dissemination services:

- The ECMWF Production Data Store (ECPDS) can be configured to disseminate the tailored data directly to the EWC storage resources to make them immediately available for processing within the cloud computing facility.

doi: 10.21957/xawe573ul9

- The Meteorological Archival and Retrieval System (MARS) stores hundreds of petabytes of meteorological and climate data which can be retrieved by users leveraging the local network access.
- The Copernicus Data Stores enable access via APIs to data from the Copernicus Climate Change and Atmosphere Monitoring Services, harnessing local network access.

Additionally, the ECMWF Aviso Data Notification service can be configured on the EWC for receiving automatic notifications of the availability of model output data and the delivery of the products from the ECPDS dissemination system, to trigger further processing workflows.

The EWC is a complementary facility to ECMWF's HPCF. The HPCF remains the most suitable facility for highly intensive CPU, I/O or interconnect batch parallel workloads. Being a managed service, the HPCF does not enable the degree of customisation and flexibility that specific use cases may require. The EWC, on the other hand, can provide a higher degree of freedom and flexibility for user-configured environments.

For example, running a public web service available on the Internet would not be suitable on the HPCF. In the EWC, users can create their own virtual environment on demand (configuring the required CPU, memory, and disk resources; assigning a public IP and domain name), deploy their web service, and finally make it available to other end users.

While the EWC and HPCF may serve different use cases, they can also be integrated as components of a more comprehensive system or application, benefiting from the respective nature of the two services. Several use cases have shown, for example, how data could be produced by a model running on the HPCF, with the output delivered to the EWC, where it could be further post-processed and served by a web visualisation service.

How can the EWC foster collaboration across the meteorological community?

The EWC is, by design, a cloud-computing-based platform built with a strong focus on European meteorological community needs. It aims to enhance collaboration across its users. The EWC platform gathers in a single place computing resources, data access, and tools to facilitate and promote data exploitation as well as the generation of value-added products and services (Figure 2).

Figure 2 The EWC provides computing resources, data access, and tools to facilitate data exploitation and the generation of value-added products and services.

The platform enables its users to collaborate on a joint virtual environment which runs in close proximity to the data. This enables the co-development of new weather and climate applications and services, and the co-design and creation of new systems and their operations.

A notable example is the EUMETNET RODEO project, whose objective is the provision of open access to public meteorological data and the establishment of a shared federated data infrastructure to develop information products and services. It is a combined effort by 11 European NMHSs, ECMWF, and the overall EUMETNET network. In the context of RODEO, the EWC is used for the development, implementation and hosting of some of the project's components. Moreover, available tools, such as the discussion platform

doi: 10.21957/xawe573ul9

(*https://confluence.ecmwf.int/display/EWCLOUDKB/EWC+Discussion+Platform*), allow users to work together, promoting information exchanges and the fundamental collaborative nature of the EWC.

The vision is to stimulate the creation of synergies and exchanges across users, to let them join forces for similar objectives, to increase reusability and reproducibility of systems and applications, and to share results and templates for common benefit.

What kind of training and outreach are offered to users?

The EWC teams at ECMWF and EUMETSAT have planned a training roadmap for the users of the platform consisting of a series of short webinars covering a wide range of topics. Their primary objective is to provide guidance to users on how to effectively use the different capabilities of the EWC. These webinars cover the basics to get started, progressively advancing to more focused thematic aspects and advanced examples. They are open to all existing and potential users of ECMWF and EUMETSAT Member and Co-operating States. All events are recorded and can be followed offline from the EWC Knowledge Base (*https://confluence.ecmwf.int/display/EWCLOUDKB/EWC+Training+and+Tutorials*).

Another important event in the calendar is the annual EWC User Workshop (*https://events.ecmwf.int/event/360/)*. Jointly organised by ECMWF and EUMETSAT, these workshops are a perfect opportunity to bring together the entire community of EWC users. The programme usually combines updates on the status of the platform and recent developments, discussions on key emerging topics, and a strong contribution from Member and Co-operating States. The latter can present their activities, inspire others, and share their experiences and feedback.

What are the main challenges and opportunities ahead?

The main challenge is to continue to ensure that the EWC offer is in line with evolving user needs. This can be achieved by providing a set of useful tools and features to keep the platform attractive and serve the meteorological community as required. To accomplish this goal, ECMWF and EUMETSAT have in place a continuous improvement process that helps to track users' requirements and to guide the development and implementation of the existing roadmap by prioritising the most critical items.

Regarding opportunities, the EWC plays a role in a wider strategy that foresees the possibility to combine the outcomes of different projects and initiatives carried out by ECMWF and EUMETSAT. The idea is to build upon these outcomes to maximise the benefits to the users of the platform. This is true especially with projects established in close cooperation with Member and Co-operating States where cloud technologies are being adopted. An example of these synergies is the EUMETNET RODEO project mentioned above and ECMWF machine learning pilot project activities.

How to get access to the service?

Eligible users may request access by contacting the ECMWF Computing Representative (*https://www.ecmwf.int/en/about/contact-us/computing-representatives*) in the corresponding Member or Co-operating State with a high-level description of their activity and an estimation of the resources required to fulfil their objectives. After assessment, if the usage is approved by the Computing Representative, the EWC support team will set up and provide access to the requested cloud resources and services.

Another way to access the EWC for users of ECMWF and EUMETSAT Member States is to apply to ECMWF Special Projects (*https://www.ecmwf.int/en/research/special-projects/special-project-application*) or the EUMETSAT Research and Development calls (*https://user.eumetsat.int/news-events/news/european-weather-cloud-research-and-development-call*) to perform research activities. This is also suitable for projects which are undertaken in cooperation between several institutions, nationally or internationally.

Any interested users, including those working internally at ECMWF and EUMESTAT, may contact the EWC support team via the EWC Support Portal (*https://support.europeanweather.cloud*) to discuss their needs and the best approach to follow.

Success stories

In the past few years, there have been a variety of interesting use cases covering a wide spectrum of technical and scientific applications. At the time of writing, about 110 different projects have made use of the EWC, in one or both sides of the cloud at ECMWF and EUMETSAT, depending on the required data sources and target applications. The increasing number of use cases confirms the growing interest from Member and Co-operating States in using and adopting the cloud technologies provided by the EWC. At the same time, the EWC is actively used internally at ECMWF and EUMETSAT within the scope of their official mandates.

doi: 10.21957/xawe573ul9

All these projects have provided invaluable feedback and insights, especially during the pilot phase. They enormously helped to improve the service and to shape its evolution, contributing to the community growth around the platform.

Data-driven machine learning applications

The increased volumes of numerical weather prediction and climate data, as well as of observations and satellite products, stimulate the research community to experiment and develop new efficient ways to exploit these data.

In this context, machine learning (ML) and artificial intelligence (AI) applications are becoming increasingly popular. This can be attributed in part to the availability of complementary computing capabilities including smaller-scale exploratory GPU environments, the abundance and accessibility of data, and finally, the availability of tools and frameworks that facilitate the development of some of these types of application. The EWC provides these elements within a coherent platform that complements the ECMWF HPCF and other resources needed for training large ML models.

Among the hosted projects, one example is the Météo-France LabIA team, which presented its activities and plans at the EWC User Workshop, held in September 2023. These plans included several research and development case studies exploiting different models and algorithms.

Another example is the transversal collaboration between Member and Co-operating States and ECMWF on a common ML pilot project to boost knowledge-sharing and target key activities, such as enhancing data-driven forecasting. As part of this work, the EWC complements the ECMWF service portfolio, particularly the HPCF, by offering collaborative environments for the participating community of experts in AI/ML.

Supporting training courses

The EWC can serve as a collaboration platform on which training courses can be run. It enables ECMWF and EUMETSAT Member and Co-operating States to conduct training sessions with a larger number of participants than would be possible in a physical classroom. For example, by using the EWC, a training course that would previously have been limited to 10 participants could have an increased capacity of 100 participants.

ECMWF has transitioned to virtual or hybrid training events, partly driven by COVID-related lockdowns. In this scenario, the EWC (including during the pilot phase) has provided support for practical and hands-on sessions through virtual classrooms. The system's flexibility means that Member and Co-operating States can also use the EWC as an infrastructure to run their own training.

EUMETSAT also runs virtual courses on the EWC platform, through which participants can easily practise using EUMETSAT data. During training courses for meteorologists, participants use virtual machines to run Jupyter notebooks, enabling them to display data from geostationary and polar-orbiting satellites quickly and easily.

The German Meteorological Service (Deutscher Wetterdienst) used the EWC during the 2020 edition of its training course on the ICON model – the German weather forecasting model. Using the EWC made the practical part of the training session go more smoothly, allowing participants to use the model to forecast the weather and put into practice what they had learned during the lectures.

International collaborations and system complementarity

The South-East European Multi-Hazard Early Warning Advisory System (SEE-MHEWS-A) was designed to improve forecasting in an area of Europe that has experienced a significant number of meteorological and hydrological hazards. Implemented by the WMO, and hosted at ECMWF during the demonstration phase, the system provides tools for forecasting such hazards and their potential impacts, on a single virtual platform.

The functionality of the EWC means that NMHSs of participating countries can deploy their own web-based applications for their users. Rather than being sent to the NMHS to be fed into their applications, the data produced on ECMWF's HPCF by Member and Co-operating States participating in SEE-MHEWS-A is sent directly into the EWC. Co-location of both services within the same data centre means that the data can be fed directly to applications efficiently and without saturating the external network. This is especially important due to the increased resolution in ECMWF's current forecasting system, Integrated Forecasting System (IFS) Cycle 48r1.

 doi: 10.21957/xawe573ul9

Support to emergency response

In March 2020, an earthquake severely damaged the headquarters of the Croatian Meteorological and Hydrological Service (DHMZ). There were no casualties and the IT infrastructure remained intact, but there was severe damage to the building hosting it. Without any prior preparation, a backup system was established on the EWC within days. Copies of the services DHMZ ran were rapidly deployed, running in parallel, providing reassurance in case of further deterioration of the building and associated infrastructure.

“I am hugely impressed by the prompt, effective response and support by ECMWF and EUMETSAT to DHMZ after the earthquake,” said Dr Branka Ivančan-Picek, former Director-General of DHMZ.

In 2021, DHMZ moved its offices and data centre to new premises. During this transition, the backup running on the EWC and ECMWF’s HPCF became operational, ensuring that DHMZ was able to provide an uninterrupted service to its users.

System interoperability

Whilst use of the EWC is primarily intended for ECMWF’s and EUMETSAT’s Member and Co-operating States, it is also available for WMO activities. The WMO Information System 2.0 (WIS 2.0) provides a framework for WMO data sharing. Users can contribute to and download data from the system.

While WIS 2.0 components could be deployed on any cloud, the EWC was selected by Météo-France, one of the main contributors to WIS 2.0, to host one of the core components during the pre-operational phase. The EWC is designed to run applications that provide a service available to external users. This example and others, showing that the EWC can be interconnected with other clouds and services around the world, is invaluable and demonstrates the outward-facing aspects of the EWC.

What is the underlying cloud computing infrastructure at ECMWF?

The EWC at ECMWF is hosted within the ECMWF Common Cloud Infrastructure (CCI). This is a cloud-computing-based IT infrastructure hosting and serving multiple projects and services offered at ECMWF. The CCI has been set up in ECMWF’s data centre in Bologna, Italy, co-located with other ECMWF facilities and services, such as the HPCF and the Data Handling System (DHS).

The CCI is also the home of the new ECMWF Copernicus Data Stores, the main data storage system and backend for the two Copernicus operational services implemented by ECMWF on behalf of the EU: the Climate Data Store (CDS) for the Copernicus Climate Change Service (C3S) and the Atmosphere Data Store (ADS) for the Copernicus Atmosphere Monitoring Service (CAMS). Figure 3 shows the relationship between the CCI, the EWC and the Copernicus Data Stores.

The cloud infrastructure is split into two production clouds, named CCI1 and CCI2. Each production cloud is hosted in a different computing hall at ECMWF’s data centre for redundancy and resilience purposes. In the current configuration, the CCI features 68 compute nodes, with a total of 19,456 cores, and over 117 TiB of memory. This infrastructure offers access to 32 Nvidia Ampere A100 80 GB graphics processing units (GPUs), required by AI and ML users and applications. An overview of CCI resources is given in Table 1.

Openstack, a well-known open-source standard cloud computing software platform, is the technology used to manage those computing capabilities and expose them to different applications, services, and users. Backing the computing capabilities, approximately 5.5 PiB of usable hard-disk-drive-backed (HDD) storage and around 300 TiB of solid-state-drive-backed (SSD) storage are available on each of the two CCI cloud clusters. Ceph is the storage solution behind them, ensuring scalable and robust access to the data. This storage capacity is used for both classic block storage by the virtual infrastructure deployed in the cloud, as well as object storage using popular APIs, such as Amazon S3 and Openstack Swift.

doi: 10.21957/xawe573ul9

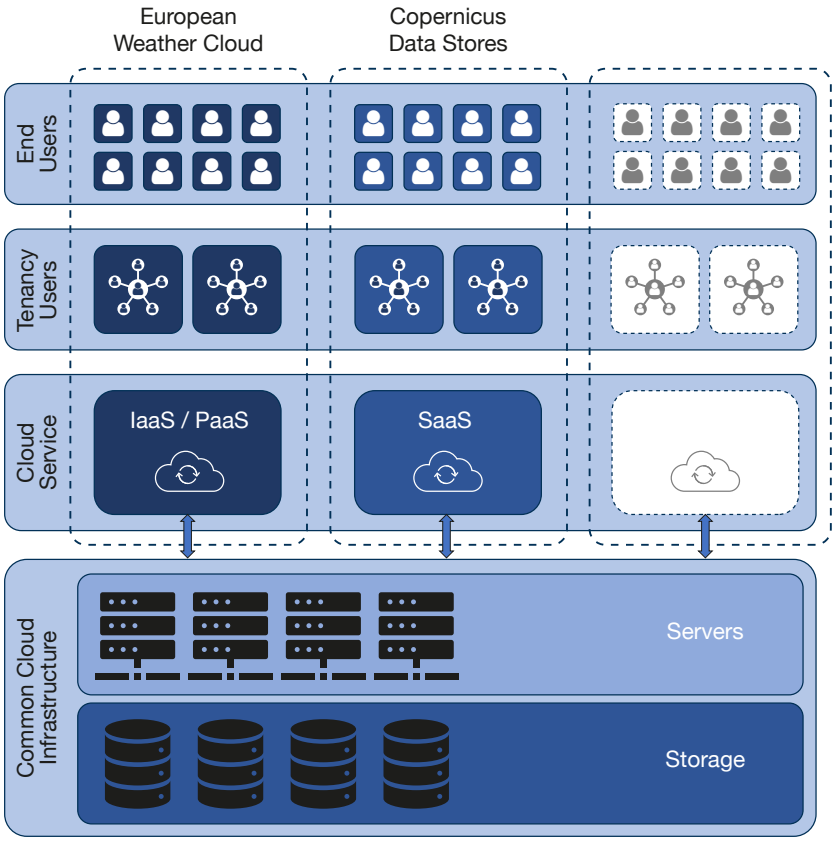

Figure 3 The CCI supports the European Weather Cloud, the Copernicus Data Stores and other applications and services. IaaS, PaaS and SaaS stand for 'infrastructure as a service', 'platform as a service', and 'software as a service'.

CCI Resources	
Compute nodes	68
Compute node central processing units (CPUs)	2 x AMD 7713 2.0 GHz, 64C/128T
Graphics processing unit (GPU) nodes	16
GPU node CPUs	2 x AMD 7543 2.8 GHz, 32C/64T
Hyper-threading cores per GPU node	128
Total GPU cards	32 x NVIDIA Ampere A100 80 GB
Total hyper-threading cores	19,456
Total physical memory	117 TiB
Total storage	11 PiB
GPU card per node	2
GPU card type	NVIDIA Ampere A100 80 GB (partitioned as virtual GPUs)

Table 1 The European Weather Cloud is one of the services hosted on the CCI. It could use approximately 50% of the overall CCI resources displayed in the table.

doi: 10.21957/xawe573ul9

Conclusion

The EWC has served many use cases and a wide range of applications in its journey from a pilot project to an operational service. The community around the platform is growing, with strong participation from existing and prospective users during relevant events, such as training and user workshops, and in the discussion platform.

The EWC continues to evolve and enrich its offer with new tools and services, improving its users' experience and further enhancing collaboration within the community.

Further reading

Pappenberger, **F.** & **M. Palkovič**, 2020: Progress towards a European Weather Cloud, *ECMWF Newsletter* **No. 165**, 24–27. *https://www.ecmwf.int/en/newsletter/165/computing/progress-towards-european-weather-cloud*

Abellan, **X.**, **K. Horvath**, **I. Pelajić** & **A. Stanešić**, 2020: Croatian met service backs up its production at ECMWF after earthquake, *ECMWF Newsletter* **No. 164**, 5–7. *https://www.ecmwf.int/en/newsletter/164/news/croatian-met-service-backs-its-production-ecmwf-after-earthquake*

Giraud, **R.**, **D. Podeur** & **T. Lacoste**, 2023: Running a Global Broker as part of the new WMO data sharing solution. *ECMWF Newsletter* **No. 177**, 32–36. *https://www.ecmwf.int/en/newsletter/177/computing/running-global-broker-part-new-wmo-data-sharing-solution*

doi: 10.21957/xawe573ul9