


ECMWF Feature article

from Newsletter Number 177 – Autumn 2023

COMPUTING

Thirty years of the ecgate
service at ECMWF



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

doi: 10.21957/umk71j84re

This article appeared in the Computing section of ECMWF Newsletter No. 177 – Autumn 2023, pp. 27–31.

Thirty years of the ecgate service at ECMWF

Paul Dando, Dominique Lucas, Umberto Modigliani

The two leading and most well-known IT systems at ECMWF are the high-performance computing facility (HPCF) and the Data Handling System (DHS). Another system, commonly known as ‘ecgate’, was available to thousands of users in ECMWF Member and Co-operating States for nearly 30 years. It was a general-purpose standalone system and the main access point to ECMWF’s Meteorological Archival and Retrieval System (MARS). It also provided computing resources enabling the post-processing of ECMWF forecast data to produce tailored products addressing specific Member and Co-operating State needs. The move of the ECMWF data centre to Bologna, Italy, and the subsequent decommissioning of the data centre at the Reading headquarters in the UK brought with it an end to the standalone ecgate system. Access to the final version of the system was closed on 15 December 2022. In this article, we look back at the history of the ecgate system, illustrate its usage, and highlight the importance of the service it provided. We also explain how this important managed service, available to Member and Co-operating State users, is now provided by part of the new HPCF in Bologna, called the ECGATE Class Service (ECS).

Key dates in the ecgate service

The ecgate service changed a lot over time. Here we present a list of significant events related to it.

1993 – First new Unix service

The first general Unix service was offered to ECMWF Member State users on 1 September 1993. This service was introduced to replace the interactive service provided on a CDC Cyber NOS/VE system, in line with the global move towards Unix systems, as users became more accustomed to the Unix environment. The first server was an SGI Challenge with two 150 MHz R4400 CPUs and was shared by Member State users and staff from ECMWF’s Operations Department. The main purpose of this first system was as a gateway for users to access services provided by other ECMWF systems. As such, it acted as a front end to the HPCF: it enabled batch job submissions, editing of files and transfer of files to and from the HPCF, Member State systems or the user archive (ECFILE). Access to MARS or ECMWF’s graphical library (Magics) was unavailable on the initial ecgate system. The system was only accessible from the national meteorological services of our Member States, and users logged in with a password.

1995 – Dedicated Member State Unix system

During the summer of 1995, a new dedicated Unix system called ‘ecgate’ was made available to Member State users. In addition to the previous service, the new system also offered a local batch service. A new Unix MARS client was now also available, enabling Member State users to download and process data from the archive directly on ecgate.

With the increased use of Unix and the Transmission Control Protocol/Internet Protocol (TCP/IP) suite, both at ECMWF and within many Member States, security aspects of remote access to the Centre’s Unix services had become a concern. On the recommendation of ECMWF’s Technical Advisory Committee (TAC), it was decided to implement a strong authentication system to control interactive login access to the ECMWF Unix systems. The solution chosen was based on the SecurID time-synchronised smart card from Security Dynamics Inc. SecurID cards were credit-card-sized, time-synchronised tokens that displayed a new, PIN-protected, unpredictable authentication code every minute, which users used instead of a password to gain login access to ECMWF Unix systems. Member State Computing Representatives were provided with some administrative control, for example, assigning spare cards or re-setting PIN codes for cards assigned to users in their country. These administrative tasks were performed on ecgate. The introduction of SecurID cards enabled ECMWF to withdraw all password-protected interactive access to the Centre’s Unix systems.

With the opening of this new Unix service to all Member State users and the introduction of the SecurID cards, we also closed direct login access to the HPCF. Member State users were therefore forced to first log in to ecgate. This change enabled us to provide independent Unix systems to Member State users and ECMWF staff.

1996–1999 – Internet access

Following ECMWF Council approval, access to ecgate via the Internet was first opened on demand at the end of December 1996. This was followed in March 1999 by the lifting of all restrictions on outgoing Internet connections from ecgate. Internet access was further extended to all Member State users in

November 2002. These moves opened access to the ECMWF computing environment to a much larger user community, including users from organisations other than national meteorological services.

2001 – Opening access to Co-operating States

In 2001, ecgate access was also granted to users from ECMWF Co-operating States. Until then, these users could only access data from the ECMWF MARS archive using a special MARS client that enabled remote access to the data. Direct access from ecgate enabled them to access and process ECMWF data more efficiently, particularly ensemble forecast (ENS) data, which was becoming more popular then. This was an example of the ‘bring the compute work close to the data’ paradigm for these users.

2002 – ECaccess service

The ECaccess software was developed by ECMWF and made available to ECMWF users. It included many functionalities to facilitate the usage of ECMWF’s computing resources. It had interactive login facilities, interactive file transfer capabilities, remote batch jobs and file handling, and unattended file transfer options. All these ECaccess services were managed on the ecgate server. Amongst others, Météo-France used the ECaccess infrastructure to enable its researchers to remotely submit experiments with its ARPEGE weather prediction model on ECMWF computing systems transparently.

2005 – Member State framework for time-critical applications

With increased user demand for running Member State operational work at ECMWF, and following approval by the ECMWF Council at the end of 2004, a framework for running Member State time-critical applications was set up at ECMWF.

- The first option enabled users to submit batch work associated with some ECaccess ‘events’. These jobs would be started when the ECMWF operational suite reached the relevant event, e.g. the 10-day high-resolution forecast (HRES) data at 12 UTC being available. Most of these jobs would run on ecgate.
- The second option enabled a Member State or consortium to run its own time-critical applications, using the ECMWF workload manager, SMS at the time. These activities were managed on ecgate.

The ECMWF shift staff monitored all these activities.

2018 – ecgate virtual machine to support ecFlow activity

HARMONIE is a mesoscale analysis and forecasting limited-area model used by several ECMWF Member States. Over the years, the HARMONIE research community using ecgate grew to over 100 users. Each user would submit one or more experiments from ecgate. By 2018, HARMONIE used the ECMWF workload manager called ecFlow, with ecFlow servers running on ecgate. To reduce interference of this intense ecFlow activity on ecgate with other ecgate activities, a new additional ecgate service was opened on a virtual machine. This machine managed only ecFlow activities, including HARMONIE experiments and some Member State time-critical applications.

2020–2022 – ecgate system freeze and final termination of service

After 2020 and until its final switch-off on 15 December 2022, ECMWF carried out only essential maintenance on ecgate so that resources could be focused on migrating the data centre to Bologna. A new service to replace ecgate, called the ECGATE Class Service (ECS), has now been implemented in Bologna as part of the Atos HPCF and provides similar functionality, while being better integrated with the HPC managed service.

A timeline showing these events is presented in Figure 1.

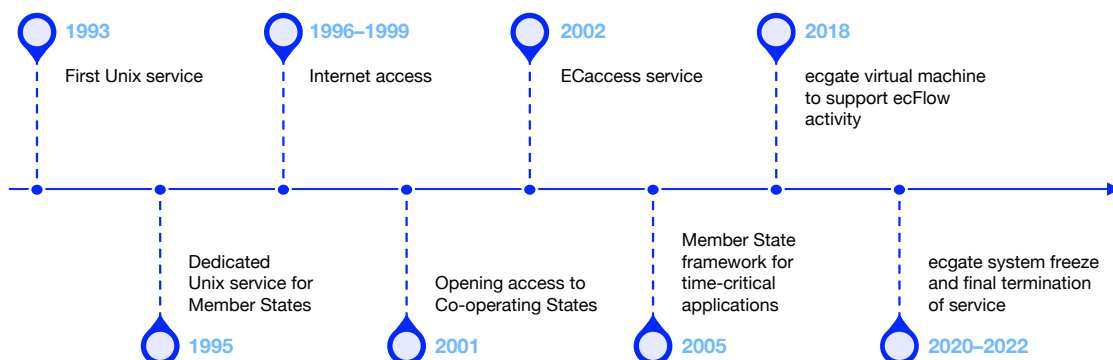


Figure 1 Timeline of the evolution of the ecgate service.

Upgrades of ecgate

The ecgate system was upgraded multiple times, as illustrated in Figure 2. These upgrades included operating system upgrades and hardware upgrades. For example, after 2013 different versions of hardware with the Linux operating system were used. Although the performance increase of ecgate broadly follows the HPC historical growth, the acquisition strategy and process for introducing new ecgate systems were independent of replacing the HPCF.




Figure 2 ecgate relative performance upgrades compared to the improvement in HPC performance.

The most crucial consideration for ecgate was that the system could handle the burst of Member State users’ activities between 05:00 and 08:00 UTC and between 17:00 and 20:00 UTC, as illustrated in Figure 3. These two activity peaks corresponded to Member State users’ time-critical option 1 jobs, accessing the latest operational HRES and ENS forecast data. These high CPU utilisations were mainly due to the MARS interpolation of real-time data. It was, therefore, essential that the ecgate system could cope with vertical or horizontal resolution increases of ECMWF’s numerical weather prediction (NWP) model, implemented as part of upgrades of the Integrated Forecasting System (IFS).

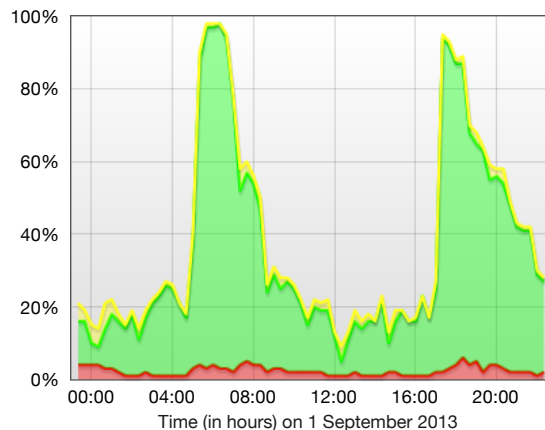


Figure 3 CPU utilisation on one node of ecgate on a day in 2013.

■ User CPU time (average of 33.64%)
■ System CPU time (average of 2.33%)
■ I/O wait time (1.93%)

A typical workday on ecgate

By 2020, we had over 3,000 registered Member and Co-operating State users with access to ecgate. Of these, over 300 users were very regularly active. The main activity on ecgate was accessing ECMWF's latest forecast data via the time-critical option 1 service. About 150 users took advantage of this service and submitted over 1,500 jobs daily.

Use of ecgate was also the most common way registered users accessed data in ECMWF's MARS archive, including historical operational data, reanalysis data and, more recently, data from the EU-funded Copernicus Atmosphere Monitoring Service (CAMS) implemented by ECMWF. These user activities would run 24/7, with a turnaround varying according to the load on the system and dependent on the load of the Data Handling System. Delays in these activities would be caught up with during quieter times, at night or during weekends.

Over a hundred users were regularly managing their own work through SMS/ecFlow on ecgate to run HARMONIE, ARPEGE, or even IFS experiments on the HPCF. In addition, many users gained access to ECMWF's graphical applications, Magics and Metview, on ecgate. Many third-party software packages were also available on ecgate and Member State users accessed them daily.

The ecgate service in hindsight

The first ecgate Unix system was set up in 1993 as a front end to the HPCF, offering only basic job and file management facilities. Over the years, ecgate was enhanced to support most user requirements, except the possibility to run computationally intensive (parallel) applications, which would run on the HPCF.

Having a separate and independent ecgate system gave ECMWF more flexibility in installing software packages needed by our users. For example, graphical packages and other Graphical User Interface applications would only be available on ecgate. The upgrade of ecgate in 2013 to Linux also gave a short-lived advantage compared to the HPCF, as it meant easier installation of third-party software packages and enabled an easier transition in 2014 to the new ECMWF HPCF, which would also run the Linux operating system.

Over the 30 years of its existence, ecgate was available to Member State users close to 100% of the time. In recent years, do-it-yourself software installation packages like conda have helped users to maintain and use their Python software stacks. On a few rare occasions, we have not been able to cater for some specialised user requirements on ecgate or, more generally, at ECMWF. This was mainly due to network and security constraints.

Figure 4 shows the ECMWF team working on the upgrade of the ecgate system in 2013, and Box A provides a view from one of our Member and Co-operating States.




Figure 4 The team who provided the ecgate upgrade in 2013.

A

“Access to the ecgate service was essential for the Hungarian Meteorological Service. Hungary has been an ECMWF Co-operating State since 1994, installing several ECMWF software packages on local servers. The right to use ecgate enabled the development of critical operational applications, such as custom clustering products for Central Europe. It also supported work that informed several master theses, such as on the dynamical downscaling of ECMWF ENS products with the ALADIN mesoscale limited-area model. These activities fostered the increasing cooperation between Hungary and ECMWF.”

Istvan Ihasz (Computing Representative and Meteorological Contact Point for the Hungarian Meteorological Service, OMSZ)

Over the years, ECMWF conducted several user satisfaction surveys. We always received very good feedback on the ease of use and reliability of the service provided by ecgate and on its importance for Member and Co-operating State users to make effective use of ECMWF data either directly or via post-processing or limited-area downscaling.

The future

A general-purpose managed computing service is now available on part of the Atos HPCF, called ECS. It offers similar functionality to that provided previously by the ecgate system and, by using the same HPCF infrastructure, it provides more efficient integration with the full HPC service, simplifying its use, maintenance, and support. We will continue offering such an important service to Member and Co-operating State users for the foreseeable future. The full Atos HPC managed service (<https://www.ecmwf.int/en/computing/our-facilities/supercomputer-facility>) complements it to run NWP models, while the European Weather Cloud (<https://www.europeanweather.cloud>) supports customised applications using ECMWF or EUMETSAT data on custom and personalised virtual platforms. More details on the ECGATE services available in Bologna can be viewed here: <https://confluence.ecmwf.int/display/UDOC/Atos+HPCF+and+ECGATE+services>. These three services complement each other and offer our users a wide range of computing platforms and advanced facilities to process ECMWF data efficiently.

© Copyright 2023

European Centre for Medium-Range Weather Forecasts, Shinfield Park, Reading, RG2 9AX, UK

The content of this document, excluding images representing individuals, is available for use under a Creative Commons Attribution 4.0 International Public License. See the terms at <https://creativecommons.org/licenses/by/4.0/>. To request permission to use images representing individuals, please contact pressoffice@ecmwf.int.

The information within this publication is given in good faith and considered to be true, but ECMWF accepts no liability for error or omission or for loss or damage arising from its use.