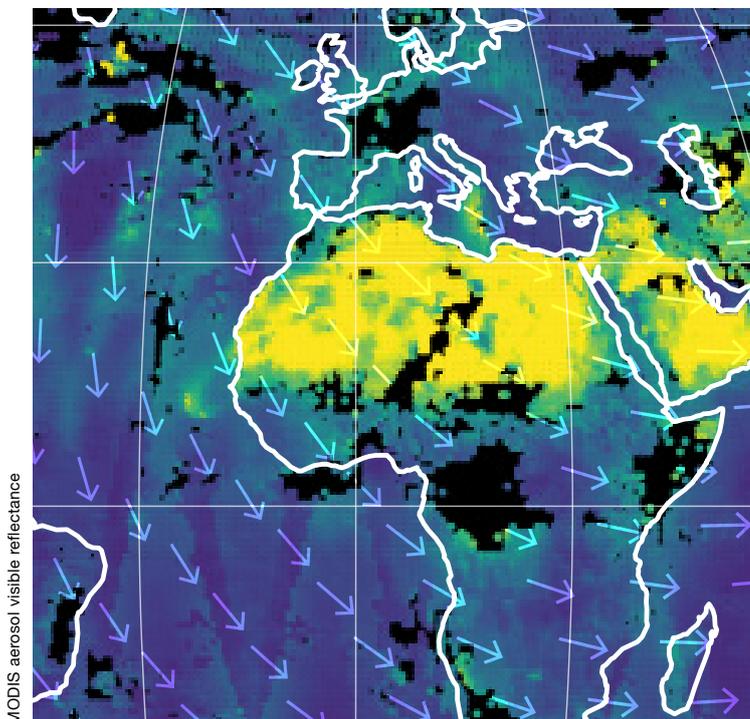




COMPUTING

GPU-adaptation of the IFS for EuroHPC machines



GPU-adaptation of the IFS for EuroHPC machines

Michael Lange, Balthasar Reuter, Ahmad Nawab, Michael Staneker, Olivier Marsden, Patrick Gillies, Zbginiew Piotrowski, Johan Ericsson, Ioan Hadade, Michael Sleight

The computational performance of the Integrated Forecasting System (IFS) has always been heavily optimised for operational use in ECMWF's data centre. Thanks to the European Commission's Destination Earth initiative (DestinE), in which the IFS forms a key part of two digital twins (DTs), the IFS is routinely run on a number of external EuroHPC systems. Many of these systems contain large numbers of graphics processing units (GPUs) to boost their computing power and support machine learning methods. In recent years, ECMWF, supported by DestinE, has worked in close collaboration with Member States, in particular Météo-France and the ACCORD consortium, with whom we share significant parts of the code, to adapt core components of the IFS forecast model to GPU accelerators. The introduction of modern software-engineering methods, paired with the thorough refactoring of the data structures underpinning the IFS, has enabled the continued optimisation for GPU accelerator architectures without impacting the central processing unit (CPU) code paths and scientific development. This has enabled ECMWF to include a GPU-enabled version of the IFS in its current high-performance computing (HPC) procurement benchmark, and it will allow for better utilisation of EuroHPC computing resources in the future.

Driving the shift to GPU computing

The use of GPUs has become ubiquitous in modern HPC. In addition to machine-learning models, and after the pioneering GPU-adaptation efforts of MeteoSwiss, several ECMWF Member States are using or exploring the use of GPU-based HPC systems for operations (Lapillonne, 2025). The ability to explore and leverage GPUs alongside CPU architectures was a driving force behind ECMWF's decision to pursue GPU support for the IFS, as it provides the ability to flexibly react to fast-changing trends in the hardware market. The desire to increase competition during HPC procurements by enabling the use of different GPU architectures from multiple vendors is a key driver behind ECMWF's GPU-adaptation strategy.

In addition, there is a long history of using external GPU resources for exploratory research runs at ECMWF that have often pushed far beyond the boundaries of operational HPC capacity (Wedi et al., 2020). With a significant increase in hardware diversity across HPC and with access to the EuroHPC

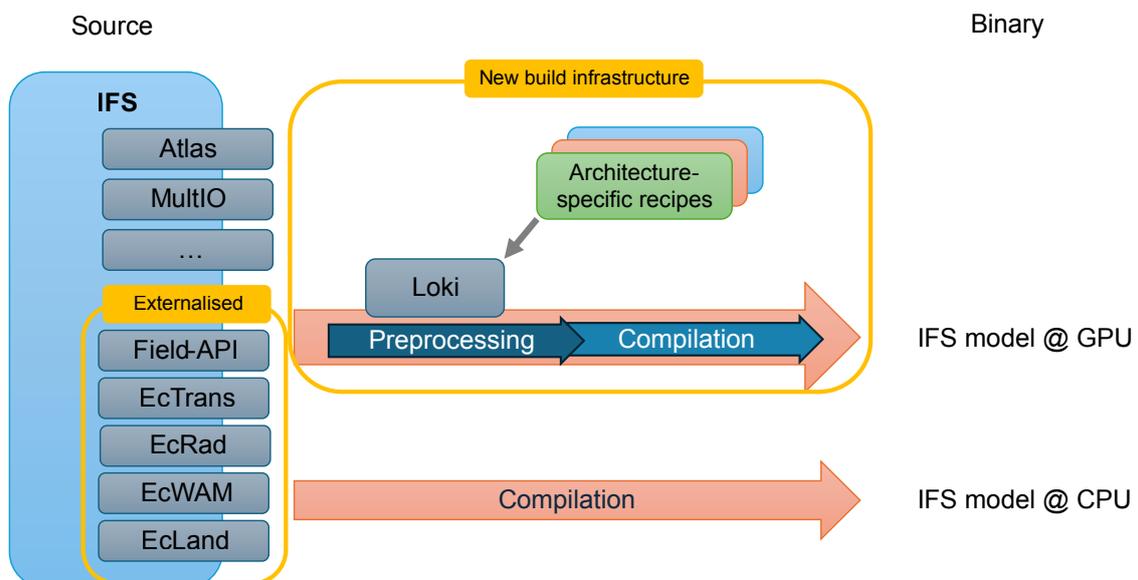


FIGURE 1 Overview of the technical infrastructure changes that enable GPU-adaptation alongside existing CPU compilation. In addition to existing external modules, like Atlas or MultiIO, new technical packages and scientific sub-models have been externalised into standalone repositories. Additional build infrastructure, including the Loki source-to-source preprocessing, is then used to automatically generate GPU-optimised code paths during the compilation process.

ecosystem, support for GPU architectures is a key requirement to fully utilise such resources to their full potential. Such resources allow ECMWF to run the IFS at resolutions beyond the in-house operational capacity, enabling research and exploration towards the next generation of physical forecast models. Recognising this trend early, ECMWF has been involved in various European projects that pioneered the use of novel software technologies on external systems (Bauer et al., 2020; Müller et al., 2019; Segura et al., 2025; Targett et al., 2021).

The EuroHPC ecosystem and Destination Earth

The EuroHPC Joint Undertaking is a major European initiative aimed at developing a world-class supercomputing ecosystem across Europe. To date, EuroHPC has procured eleven state-of-the-art supercomputers, four of which have entered the top ten in the Top500 list of the fastest supercomputers (<https://www.top500.org/>), including the first European exascale supercomputer, JUPITER (<https://jupiter.fz-juelich.de/>).

Several of the supercomputer systems procured by EuroHPC are the target platforms for the DestinE digital twins (Geenen et al., 2024; Wedi et al., 2025). The IFS underpins the Continuous Global Extremes Digital Twin (Extremes DT) and is used in the Climate Change Adaptation Digital Twin (Climate DT; Doblas-Reyes et al., 2025), coupled to the NEMO 4 or FESOM 2 ocean model. Consequently, the IFS must be adapted and tuned to run optimally on these EuroHPC systems. ECMWF develops and operates the Extremes DT and collaborates with the Climate DT consortium on porting and optimising the IFS-based Climate DT configurations for the target platforms.

Many EuroHPC systems feature two main partitions of different node types:

- A traditional CPU-only cluster consisting of nodes with two CPUs each (e.g. LUMI-C, LEONARDO-DCGP, MareNostrum5-GPP), connected with a high-bandwidth network.
- An accelerated partition comprised of nodes with one or two CPUs and four GPUs each (e.g. LUMI-G, LEONARDO-BOOSTER, MareNostrum5-ACC).

The new JUPITER system marks a shift in this architecture: it currently consists solely of an accelerated partition (JUPITER-BOOSTER), with a much smaller CPU-only cluster planned to be installed later. Moreover, the four GPUs per node are complemented by the same number of CPUs, with each pair working in tandem (a so-called “super-chip”). This allows for faster data copies between CPU and GPU memory spaces and even a unified memory space that removes the need for explicit data transfers in code.

Software architecture

Contrary to conventional GPU porting efforts, where technical work is performed on a fixed version of the code, the IFS porting strategy focused on creating a sustainable solution in which scientific development can continue alongside the technical adaptation to GPUs without compromising operational performance on the CPU-based in-house HPC system. In close collaboration with Member States, a set of modern software engineering methods was devised that allowed a “GPU mode” to be embedded in the IFS without impeding the CPU code paths. The capability to run the IFS on GPUs and CPUs from the same code base allows the continual adaptation, testing and maintenance of both code paths, as well as future extensions to support multiple programming models and hardware vendors.

To achieve this separation of concerns, in which technical optimisations are done independently of scientific changes, a combination of library development, data-structure refactoring and source-to-source translation tools was used to create GPU-specific build modes for the Fortran source code. A particular emphasis was put on the extraction of several model sub-components into standalone open-source software libraries. This allowed ECMWF to benefit greatly from ongoing collaborations with HPC hardware vendors that provided significant performance optimisations to core components of the IFS. The provision of freely accessible standalone mini-apps, the IFS “dwarfs” (Müller et al., 2019), enabled not only the integration of hardware-specific optimisations, but also significantly increased the test coverage for technical changes before integration into the full system. Notable standalone packages include the spectral transform library ecTrans (<https://github.com/ecmwf-ifs/ectrans>), the wave model ecWAM (<https://github.com/ecmwf-ifs/ecwam>) and the radiation model ecRad (<https://github.com/ecmwf-ifs/ecrad>).

A pivotal element in GPU porting is the management of program data across two separate memory spaces: GPU memory and CPU memory. This challenge is addressed through the “Field-API” (https://github.com/ecmwf-ifs/field_api), an array data management abstraction that handles data movement between CPU and GPU while remaining compatible with the complex memory data layout in the IFS.

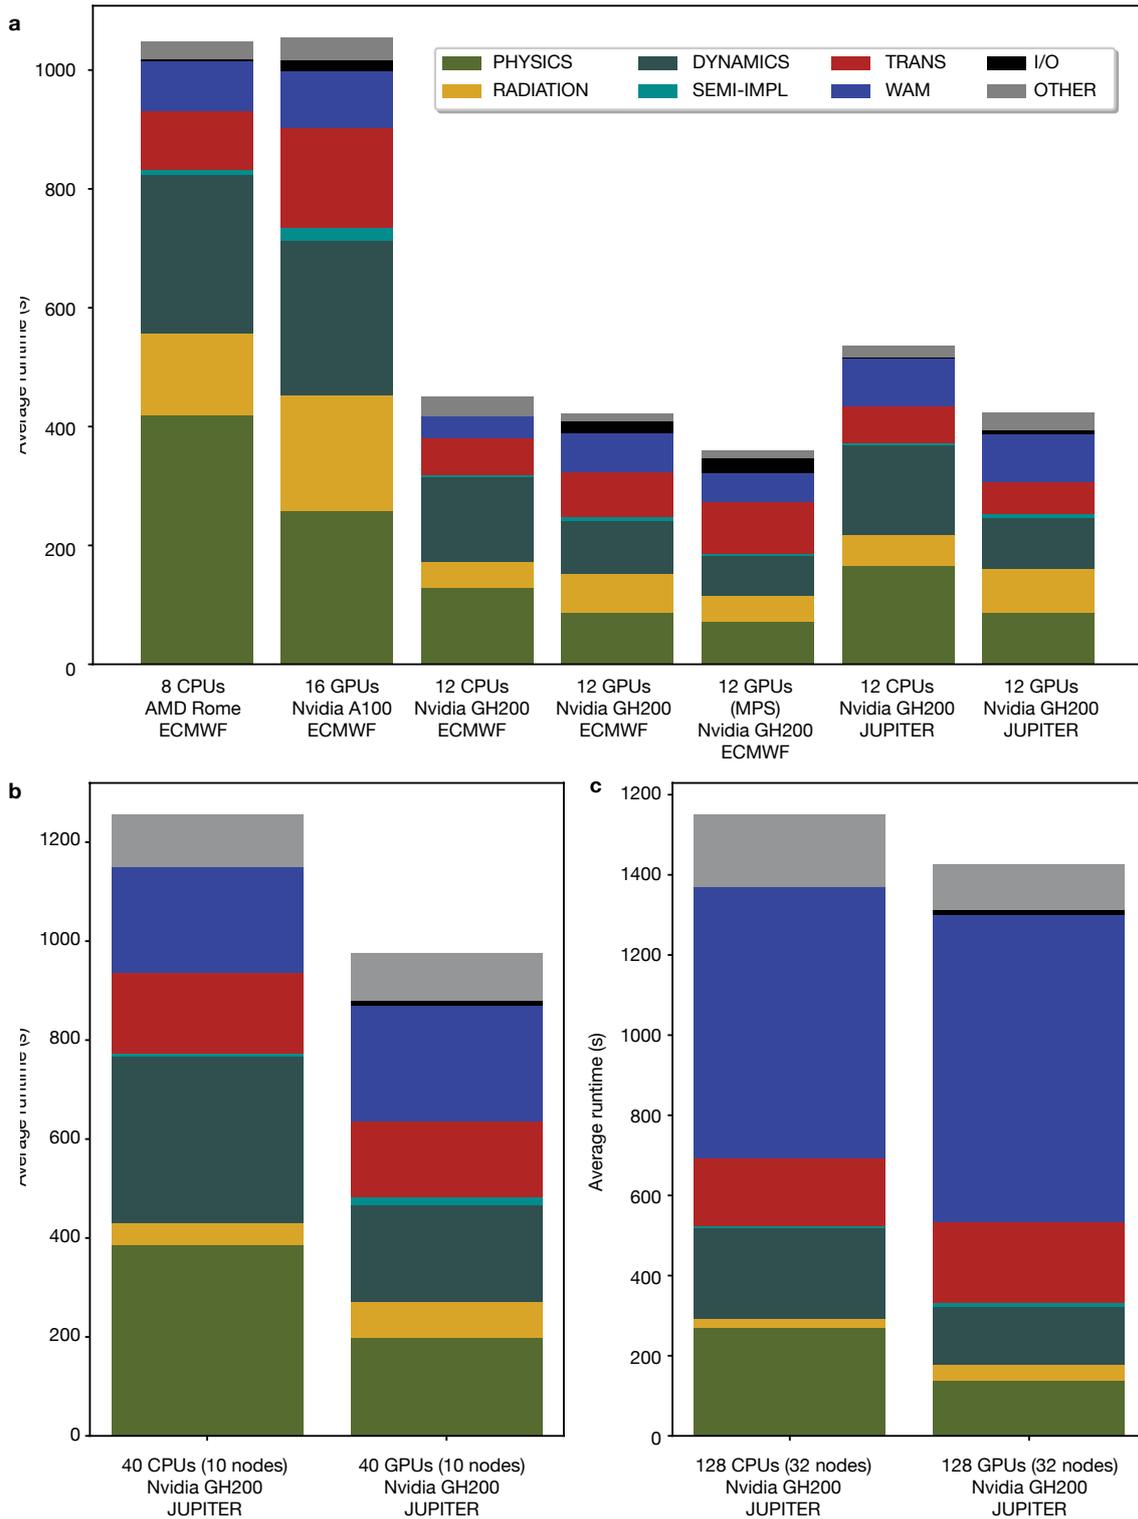


FIGURE 2 CPU–GPU performance comparison across resolutions and architectures. (a) Runtime breakdown of 48-hour forecast at 18 km resolution (TCo639) on GPU partitions of the ECMWF HPC and JUPITER. The initial column shows the runtime breakdown per component on the current operational AMD “Rome” CPU architecture and compares it to discrete GPUs (Nvidia A100). The performance of next-generation CPU-GPU “superchips” (Nvidia GH200, “Grace-Hopper”) is then shown for ECMWF’s internal machine learning partition and the EuroHPC system JUPITER. The use of GPU-oversubscription via Nvidia’s Multi-Process Service (MPS) on the internal system brings further performance gains for GH200 but was not yet available on JUPITER. (b) Comparison of CPU and GPU runtimes for 48-hour forecast at 9 km resolution (TCo1279) on JUPITER, showing per-component breakdowns on minimal node counts (10 nodes, 40 CPUs/40 GPUs). (c) Comparison of CPU and GPU runtimes for 12-hour forecasts at 4.4 km resolution (TCo2559) on JUPITER, using 32 nodes (128 CPUs/128 GPUs).

This dedicated Fortran infrastructure library provides vendor-agnostic support for GPU data offload and delivers optimised backends that use low-level programming paradigms, such as CUDA and HIP, to ensure efficient data movement and layout in GPU device memory.

The key to separating the scientific algorithm from the necessary code optimisations is the source-to-source translation tool Loki (<https://github.com/ecmwf-ifs/loki>), a vital component of the Digital Twin Engine (DTE) developed at ECMWF with support from DestinE. Loki is a Python package that allows HPC specialists to encode a set of complex source code transformations as automated “recipes” that are applied during the model compilation process to transform and optimise the CPU code to GPUs. These recipes are highly bespoke and use IFS coding conventions to derive GPU-optimised code from conventional IFS subroutines, using domain-specific knowledge of the underlying algorithms.

Implementing multi-architecture capabilities

The extraction of several key abstractions and scientific sub-components into open-source, standalone modules was key to implementing the IFS’s new multi-architecture capabilities. This approach not only allows the development of prototypes of new technical features independently but also enables external experts to optimise performance for specific hardware. This ability was crucial, as the IFS includes a range of diverse computational patterns, often requiring tailored porting strategies for CPU and GPU. For example, the spectral transforms, central to the model’s dynamical core, demand highly customised optimisations to scale on large HPC systems. While these optimisations often differ between CPU and GPU, the externalised ecTrans package now offers dedicated code paths for both, unified under a common library interface. Thanks to its open-source release, the GPU paths have also been optimised by experts from Nvidia and AMD for both GPU architectures, resulting in a highly optimised library that supports a variety of CPU and GPU architectures.

As highlighted in Figure 1, for model components where code replication was avoided through automated source-to-source translation, extracting small, representative examples also proved essential for sustainable multi-architecture support. This included both individual routines – like the CLOUDSC cloud microphysics scheme – and entire sub-models such as ecWAM (waves) and ecRad (radiation). These examples enabled the development of tailored solutions combining GPU-enabled data structures, automated translation via Loki, and manual porting to test GPU support. Many of these efforts were highly collaborative, involving Member States, European project partners, the DestinE initiative, and industry experts. This open development model supports ongoing improvements to translation recipes and GPU-specific optimisations, allowing ECMWF to react quickly to emerging hardware trends.

Sustainability and testing

ECMWF has recently added a new GPU partition to its Bologna HPC facility. In addition to enabling and enhancing the machine learning abilities at ECMWF (Pappenberger et al., 2024), this new hardware allows the continuous-integration testing (CI) of the newly developed GPU capabilities. The automated deployment of GPU-enabled test runs, in conjunction with the automated generation of GPU-enabled code paths from standardised CPU code, allows the correctness of the GPU code paths to be routinely checked and verified with only a moderate maintenance overhead.

Thanks to the access provided by DestinE, this capability is even extended to external HPC systems (e.g. LUMI) for individual externalised model components, such as the spectral transform library ecTrans. As a result, the IFS has become a true multi-architecture model that is tested and deployable across a range of HPC systems. The sustainability of this approach – and the effectiveness of separation of concerns – was demonstrated when the GPU-adaptation recipes developed for IFS Cycle 49r1 were migrated successfully to the next version, Cycle 50r1. This was achieved in a relatively short time by a few technical experts, without the need for any intervention from scientific developers.

Status and initial results

With the recent rise of GPU architectures to the forefront of HPC systems, GPU hardware characteristics have become increasingly diverse. In addition to traditional discrete GPU systems, where data transfers between CPU and GPU memory regions are often the main performance bottleneck, shared-memory systems have been developed. These combine CPU and GPU architectures in the same chip or contain dedicated hardware for more efficient CPU–GPU memory transfers. The ECMWF data centre contains two GPU partitions, representing both hardware types.

ECMWF’s overall GPU-adaptation strategy aims to stay agnostic to specific GPU types. The recent focus has been on porting most of the atmospheric time step of the IFS to GPU to allow data to stay resident in the GPU memory, reducing the number of data transfers required. This development allows

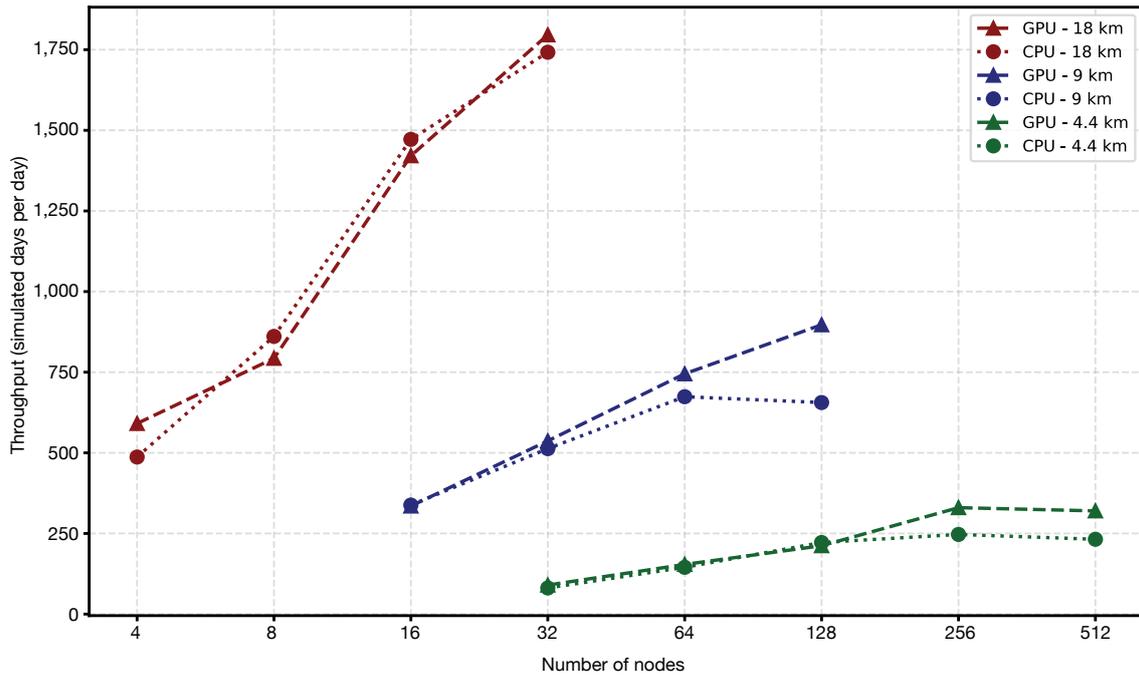


FIGURE 3 Strong scaling of CPU and GPU runs on JUPITER for research and operational resolutions. This highlights that initial GPU optimisations are already sufficient to preserve the excellent strong scaling efficiency of the original IFS implementation for CPUs.

the systematic testing and assessment of performance of the GPU port and enables the verification that operational CPU performance has not been affected. The most recent GPU-adaptation efforts have also been aligned with the latest scientific developments in IFS Cycle 50r1 – a step made possible by the high degree of automation involved in the derivation of the GPU code paths. The resulting GPU capabilities are intended to be sustainable across several cycles, allowing further optimisation in ongoing and future developments.

Initial results are already showing promise. A recent performance comparison is shown in Figure 2, where discrete GPUs (Nvidia A100) in the ECMWF HPC centre are compared with next-generation shared-memory GPUs (GH200, “Grace-Hopper”), in both the new in-house GPU partition and the EuroHPC system JUPITER. The initial, *non-optimised* GPU performance is already competitive with the *highly optimised* CPU performance, both at development and operational resolutions. Considering the computational power of the next-generation “Grace” CPUs used in the respective systems and the known optimisation headroom for the GPU implementation, these results provide a strong starting point for further improvements.

Importantly, as shown in Figure 3, the excellent multi-node scalability of the IFS is also preserved on GPUs, suggesting that future explorations of higher resolutions beyond operational constraints will be able to utilise GPU-driven machines just as effectively as they will CPU-based systems. Given the prevalence of GPU systems in the modern HPC ecosystem, this ability to compare and utilise different types of high-end systems with the IFS is essential to objectively assessing and adapting to changes in the HPC landscape over the coming years.

Conclusion and next steps

The initial adaptation of the IFS to GPU accelerators has been achieved and concluded with the inclusion of the novel GPU build modes in the latest HPC procurement benchmark release. Through the involvement in external projects and the DestinE initiative, ECMWF will continue its efforts to maintain and improve upon these capabilities to best utilise any available internal or external computing resources. The high degree of automation and the enhanced automated testing capabilities on internal GPU hardware enable these features to be maintained sustainably. As a result, the novel GPU capabilities are planned to enter the mainline code in IFS Cycle 50r2 and remain available to ECMWF researchers. The capabilities will be retained in subsequent IFS releases, allowing potential inclusion in future releases of OpenIFS.

In addition to novel technical capabilities, one of the key outcomes of the GPU-adaptation efforts has been the restructuring of several components of the forecast model and the improvement in technical testing infrastructure. ECMWF recognises the need to further address existing technical debt in the IFS code base and aims to improve the sustainability of the physical model, ensuring that future technical changes can be implemented with the required agility. The FORGE (Forecast System Regeneration) project (Sleigh et al., 2025) has been established to spearhead the modernisation of the forecast system and associated software infrastructure. This initiative will work closely with Météo-France and ACCORD partners to ensure that ECMWF continues to adapt to the increasing technical demands in a highly diverse environment.

Further reading

Bauer, P., T. Quintino, N. Wedi, A. Bonanni, M. Chrust, W. Deconinck et al., 2020: The ECMWF Scalability Programme: progress and plans. *ECMWF Technical Memorandum No. 857*. <https://doi.org/10.21957/gdit22ulm>

Doblas-Reyes, F. J. et al., 2025: The Destination Earth digital twin for climate change adaptation. *GMD [preprint]*. <https://doi.org/10.5194/egusphere-2025-2198>

Geenen, T. et al., 2024: Digital twins, the journey of an operational weather prediction system into the heart of Destination Earth. *Procedia Computer Science*, **240**, 99–109. <https://doi.org/10.1016/j.procs.2024.07.013>

Lapillonne, X., et al., 2025: Operational numerical weather prediction with ICON on GPUs (version 2024.10). *Geoscientific Model Development [preprint]*. <https://doi.org/10.5194/egusphere-2025-3585>

Müller, A. et al., 2019: The ESCAPE project: Energy-efficient Scalable Algorithms for Weather Prediction at Exascale. *Geoscientific Model Development*, **12(10)**, 4425–4441. <https://doi.org/10.5194/gmd-12-4425-2019>

Pappenberger, F. et al., 2024: Artificial intelligence and machine learning: revolutionizing weather forecasting. UN United in Science 2024 Report.

Segura, H. et al., 2025: nextGEMS: entering the era of kilometer-scale Earth system modelling. *Geoscientific Model Development*, **18(20)**, 7735–7761. <https://doi.org/10.5194/gmd-18-7735-2025>

Sleigh, M. et al., 2025: Modernisation of the Integrated Forecasting System. *ECMWF Newsletter No. 182*, 19–23. <https://doi.org/10.21957/m9ad5hv72s>

Targett, J. et al., 2021: Systematically migrating an operational microphysics parameterisation to FPGA technology. *2021 IEEE 29th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pp. 69–77. <https://doi.org/10.1109/FCCM51124.2021.00016>

Wedi, N. et al., 2020: A baseline for global weather and climate simulations at 1 km resolution. *Journal of Advances in Modeling Earth Systems*, **12(11)**. <https://doi.org/10.1029/2020MS002192>

Wedi, N. et al., 2025: Implementing digital twin technology of the earth system in Destination Earth. *Journal of the European Meteorological Society*, Volume 3, 100015. <https://doi.org/10.1016/j.jemets.2025.100015>

© Copyright 2026

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.