Redefining Weather Forecasting Systems: The Transition to ICON and Alps

Mauro Bianco*
mauro.bianco@cscs.ch
ETH Zurich, Swiss National Supercomputing Centre (CSCS)
Lugano, Switzerland

Matthias Kraushaar
matthias.kraushaar@cscs.ch
ETH Zurich, Swiss National Supercomputing Centre (CSCS)
Lugano, Switzerland

Roberto Aielli
roberto.aielli@cscs.ch
ETH Zurich, Swiss National Supercomputing Centre (CSCS)
Lugano, Switzerland

Oliver Fuhrer
Federal Office of Meteorology and Climatology MeteoSwiss
Zurich, Switzerland
ETH Zurich
Zurich, Switzerland

Thomas C. Schulthess
thomas.schulthess@cscs.ch
ETH Zurich, Swiss National Supercomputing Centre (CSCS)
Lugano, Switzerland

ABSTRACT

The transition of MeteoSwiss operational weather forecasting from COSMO to the ICON model represents a major modernization in meteorological services, integrating software-defined infrastructure to improve flexibility, scalability, and resilience. The migration also involved significant hardware upgrades, from fixed systems with K80 GPUs to flexible architectures using V100 and later A100 GPUs, supported by *Alps* infrastructure developed by CSCS that is based on the Cray EX product line of HPE.

The evolution included consolidating previously discrete systems into a shared cloud-like infrastructure with high-speed networking and advanced resource management using Cray System Management (CSM) framework and versatile software defined cluster (vCluster) along with User Environments (uenv) technologies. This approach allows for the deployment of isolated and customizable configurations for various services on the infrastructure by minimizing the overall complexity of the system. Extending Alps to the data center facility at EPFL in Lausanne, improved resiliency, enabling facility updates and maintenance without disruption of service and supporting critical applications like weather forecasting, disaster response, and research.

The transition not only upgraded technical capabilities, but also established a foundation for scalable, future-proof meteorological services. This evolution ensures that MeteoSwiss can meet the growing demands of modern weather and climate science.

*This work reflects the collective contributions of CSCS and MeteoSwiss teams, who have been essential to make possible the contextual transition to a new model and new infrastructure.

CUG25, May 04–08, 2025, New Jersey

ACM ISBN 978-1-4503-XXXX-X/2025/05
https://doi.org/XXXXXXX.XXXXXXX

CCS CONCEPTS

• **Information systems → Computing platforms**; • **Computer systems organization → Cloud computing**; **Heterogeneous (hybrid) systems**; • **Computing methodologies →** *Massively parallel and high-performance simulations.*

KEYWORDS

Weather Forecasting, HPC,

ACM Reference Format:
Mauro Bianco, Matthias Kraushaar, Roberto Aielli, Oliver Fuhrer, and Thomas C. Schulthess. 2025. Redefining Weather Forecasting Systems: The Transition to ICON and Alps. In *Proceedings of Computing Horizons: CUG 2025 (CUG25).* ACM, New York, NY, USA, 9 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Accurate and timely weather forecasting is essential for informed decision making in sectors such as aviation, agriculture, and disaster management. To address increasing demands for precision and adaptability, MeteoSwiss adopted the Icosahedral Nonhydrostatic Weather and Climate Model (ICON) [4, 18], originally developed at the Deutscher Wetterdienst (DWD). Simultaneously, the introduction of the Alps infrastructure - based on HPE's Cray EX product line and engineered for flexibility, scalability, elasticity, and resilience - enabled MeteoSwiss to transition from traditional, rigidly defined High-Performance Computing (HPC) systems to a dynamic, software-defined framework.

Previously, MeteoSwiss relied on COSMO-NExT and COSMO-Modinterim [15, 3], models that were implemented on GPU accelerated architectures [7] using discrete, dedicated vertically integrated systems. Such architectures imposed fixed resource distributions, limiting adaptability to evolving computational demands and the inherent physical characteristics of the hardware. This fixed allocation had two major consequences: first, ensuring service availability required substantial redundancy, thereby increasing cost; and second, any improvements in forecasting performance resulted in suboptimal resource utilization. The inflexible nature of these systems meant that accommodating new or enhanced requirements necessitated a complete system overhaul.

The convergence of HPC and cloud infrastructures - driven by innovative initiatives such as those at the Swiss National Supercomputing Centre (CSCS) - offered a pathway to overcome these limitations. By enabling MeteoSwiss services to run on a larger shared infrastructure rather than on isolated, dedicated systems, many drawbacks of the closed-system approach can be mitigated, paving the way for more flexible resource management and adaptive forecasting capabilities.

This convergence spans all layers of the system architecture, from hardware to software. At the hardware level, the ability to isolate components such as compute nodes, network traffic, and storage is crucial for ensuring system independence and improving reliability. Although certain features - like full storage multi-tenancy - are still under development, the overall direction is clear. Building on this foundation, CSCS has introduced vCluster [11] technology, which enables the flexible deployment of system-level software and services on dedicated computing resources for tenants such as MeteoSwiss. Furthermore, atop the vCluster framework, the dynamic software environment deployment system, uenv [5], facilitates the isolation of application environments. This approach allows services and applications to execute independently of a fixed setup, overcoming a constraint that has long limited traditional HPC systems.

This modernization also sets the stage for more flexible integration of new services and workflows, such as tighter interoperability with machine learning techniques to further enhance forecasting capabilities. The technology can also streamline hardware upgrades: the recent migration from NVIDIA V100 to A100 GPUs - transitioning from a dedicated system to the shared Alps infrastructure - illustrates how future upgrades could be more readily tested and deployed by migrating portions of the infrastructure using the same set of tools across platforms. Additionally, the capabilities provided by the Cray EX Alps platform, combined with vClusters and uenvs, have enabled the establishment of a geo-distributed infrastructure at EPFL, further enhancing the resilience of MeteoSwiss services and ensuring continuity even during maintenance or unexpected disruptions at a single site.

This paper explores the motivations, technical challenges, and key innovations behind this transformation, providing insights into how MeteoSwiss is redefining operational weather forecasting.

2 ENHANCING WEATHER FORECASTING CAPABILITIES

2.1 Improved Forecasting Accuracy

The approach to producing numerical weather forecasts remains largely consistent when transitioning from COSMO to ICON as both models solve the governing partial differential equations. However, several key differences give ICON distinct advantages:

- **Icosahedral Grid:** As indicated by its name, ICON employs an icosahedral grid composed of triangular elements. Unlike COSMO's regular latitude-longitude grid, this design mitigates grid distortion - particularly near the polar regions - resulting in a more accurate representation of the Earth's surface and reducing grid-imprinted biases.
- **Unstructured Grids for Better Topography Representation:** ICON's irregular grid allows for enhanced resolution of complex topographical features. This is particularly beneficial in regions with steep orography, such as the Alpine area that dominates much of Switzerland. For example, Figure 1 illustrates how the Swiss peak Säntis is more accurately rendered using the triangular grid (right) compared to a regular Cartesian grid (left).
- **Advanced Numerical Methods and Grid Refinement:** ICON's numerical techniques are better suited to handle the steep mountain slopes captured by its grid. Additionally, the model supports localized grid refinement, which allows for higher spatial resolution in areas of particular interest.

The points highlighted above demonstrate that MeteoSwiss benefits from a more sophisticated forecasting model. ICON enables predictions at finer resolutions with improved accuracy. Notably, its capacity to dynamically adapt grid resolution enhances computational efficiency and reduces operational costs.

Another significant motivation for adopting ICON is its support by a broader community, especially as COSMO approaches the end of its development lifecycle. These advantages collectively made the transition to ICON the most rational choice for advancing MeteoSwiss's numerical weather forecasting capabilities.

2.2 Supporting Disaster Response and Real-Time Analytics

CSCS developed the concept of "a versatile research infrastructure" that is implemented on the Alps infrastructure [12]. This is the base abstraction upon which the services offered by CSCS are and will be build.

The flexibility offered by the new infrastructure significantly enhances real-time weather analysis and disaster response capabilities. By dynamically allocating computing resources, MeteoSwiss can prioritize critical forecasting operations during extreme weather events, ensuring timely and accurate predictions. With the adoption of the software-defined platform discussed in the following sections, this capability can be further leveraged to provide improved responses, as computing resources can be provisioned on demand to supply additional power when needed.

The same Alps infrastructure also powers other platforms, such as: the User Lab, which provides traditional HPC services; the Merlin 7 system of the Paul Scherrer Institute (PSI), delivered by CSCS as infrastructure as a service [10]; the Machine Learning (ML) platform, supporting large language models (LLMs) and ML for scientific workloads [14]; and, most relevant to this paper, the Climate and Weather platform developed within the EXCLAIM project [16]. The latter also targets the ICON model and is designed to support both high-resolution climate simulations and data, as well as lower-demand climate and weather science workloads. Figure 2 depicts the concept of the partition of the Alps infrastructure in different platforms.

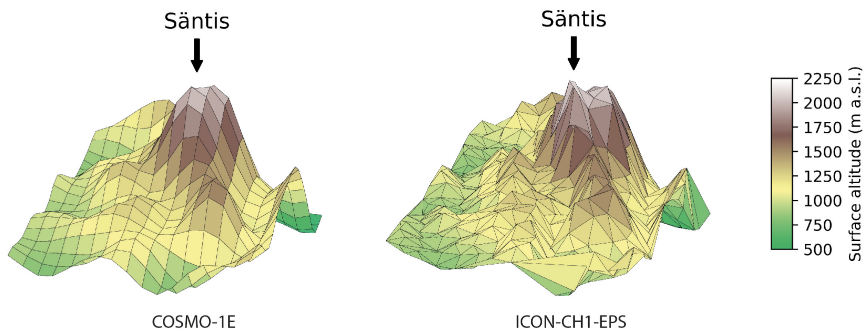

Figure 1: Comparison of the regular grid of COSMO (left) and the icosahedral grid of ICON (right) for Säntis, an alpine peak in northeastern Switzerland. Both grids show the representation at a horizontal resolution of 1km. (Source MeteoSwiss)

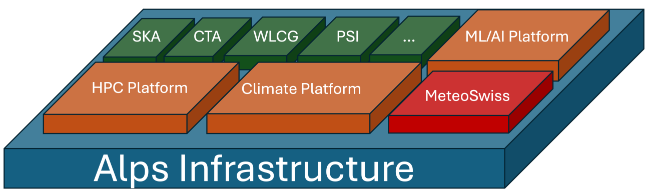

Figure 2: Illustration of different platforms on top of the Alps infrastructures. Each platform is composed of at least one vCluster, but can also consist of multiple vClusters. MeteoSWiss vCLuster sits on a portion of the infrastructure under UPS.

3 SOFTWARE-DEFINED INFRASTRUCTURE: A FLEXIBLE AND SCALABLE APPROACH

MeteoSwiss's transition to ICON is accompanied by a concurrent migration to the Alps infrastructure. Alps introduced a new architecture that represents a shift from traditional HPC systems to a convergence of HPC and cloud services.

Traditional HPC systems are typically vertically integrated, providing a unified interface to all users as shown in Figure 3. In contrast, the novel cloud-oriented approach offers tailored interfaces for different user groups. With this approach, we no longer refer to Alps as an "HPC system" but rather as a "research infrastructure" that delivers essential resources such as compute nodes, storage, connectivity, together with a software stack such a base operating system. This infrastructure serves as the foundation upon which platforms and services can be transparently deployed. The platforms built on top of this infrastructure can define various access methods, computing capabilities, queuing strategies, storage options, and more.

To enable the deployment of various platforms on this infrastructure, CSCS developed two key software technologies. The first, vClusters [6, 1, 11], allows flexible deployment of HPC systems on the infrastructure. The second, User Environments (or `uenvs`) [5], provides a customized software stack atop the vClusters. The next two sections will briefly describe the main characteristics of these important components.

3.1 Versatile Software-Defined Cluster

A vCluster is a logical partition of the supercomputing resources available within the infrastructure where platform services are deployed. It serves as a dedicated environment tailored to support a specific platform. The allocation of resources and services for each vCluster is defined by a set of configuration files that is fed into automated deployment pipelines. Once deployed, the vCluster becomes immutable, a design choice that prevents the environment from becoming unmanageable over time due to untracked, incremental changes. Instead of documenting alterations on a live system, any modifications to the deployment recipes are tracked and versioned using `Git` repositories. The recipes themselves serve as the primary documentation: once specified, they are rigorously tested, and only upon successful validation is the deployment pushed into production. In case of problems, the previous recipe can be redeployed and new tests can be conducted in isolation.

For example, vClusters are composed of *vServices* such as the job scheduler Slurm [17]. The deployment recipe for a vCluster delineates which components and versions of the service are allocated to the service plane of the infrastructure such as a Slurm control service and database while assigning Slurm daemons to the compute nodes. This design allows each platform on the Alps infrastructure to run its own instance of Slurm, with policies that reflect the specific agreements between platform users and owners.

3.2 User Environments

Similarly to vClusters, user environments (`uenvs` [5]) are built on the principle of immutability. They are defined in YAML files as declarative recipes, which are then compiled into images stored in `squashfs` files. These images are made available in a repository, allowing users to fetch and deploy them within their own shell environments. When activated, a user environment provides all the necessary facilities to run - and, if needed, build - specific

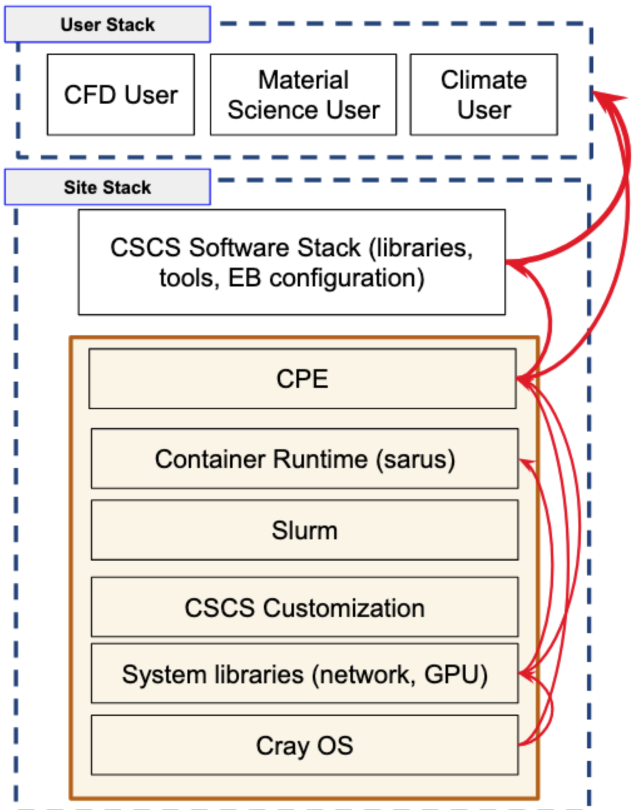

Figure 3: The "standard" HPE-EX software stack, featuring the Cray OS, drivers, CPE, and site-specific software integrated into the system image, with site-provided software installed on a shared file system. User-installed software depends on these underlying layers. The red arrows illustrate how changes in one layer can cascade, necessitating rebuilds or reconfigurations in other layers. [5]

applications. Switching between different user environments is efficient, requiring only a simple read-only `mount` operation on an image and the creation of a new shell environment, which can be easily reversed.

A user environment provided by `uenvs` can range from a standalone application to a rich ecosystem organized in `modules`. This modularity allows platforms that rely on the traditional organization of legacy supercomputers to continue operating seamlessly with `uenvs`. It facilitates a smooth transition for users from a conventional, fixed environment to a more dynamic one while preserving compatibility with established workflows.

One key advantage of `uenvs` compared to a vertically integrated user environment is that each platform can independently define and design its user environment. This approach minimizes the complexity of dependency tracking as it eliminates the need to accommodate a vast array of use cases within a single system. Furthermore, it empowers platform managers and owners to develop, test, and deploy their own `uenvs` in coordination with the vCluster technology - often without the direct involvement of central system engineers. As a result, the number of services and platforms that can be supported scales far beyond what is typically feasible in a traditional supercomputing facility, where a single monolithic resource is shared among all users.

Figures 3 and 4 illustrate the differences between the two approaches. In a traditional setting (Figure 3), the software stack is provided by the site and all dependencies are managed at the system level. In contrast, the vCluster+`uenv`-based setup allows for the development and testing of different environments in isolation. Every change must be evaluated separately, and because the images are immutable, modifications to the environment can only occur after the complete cycle of building and testing is successfully completed prior to deployment.

3.3 Evolution of the Weather Forecasting System

Historically, MeteoSwiss operated dedicated HPC clusters with static resource allocations. In the early phase (2016-2019), the organization deployed twin systems with a fixed resource distribution between operational and R&D services, utilizing NVIDIA K80 GPUs. This configuration made MeteoSwiss the first national weather service worldwide to compute forecasts using a GPU-based system in operational settings [7, 8].

A subsequent evolution (2019-2024) introduced a single redundant high-speed network domain via NVIDIA's Bright Cluster management system. Rather than maintaining two separate systems, with one dedicated to R&D and acting as a fail-over, the V100-based system consolidated these roles into a single environment capable of partitioning computing resources dynamically according to demand. This development represented an important first step towards establishing a flexible software-defined infrastructure [2].

The transition to ICON has further advanced this approach by utilizing A100 GPUs alongside the software-defined methodologies described in previous sections. Since MeteoSwiss began its deployment on the Alps infrastructure during its early days, some solutions remain in place that which have yet to fully leverage the capabilities of vClusters and `uenvs`. The stringent requirements of MeteoSwiss operations have contributed to a slower pace of final adoption compared to more experimental settings. Nonetheless, the path towards full integration of software-defined platforms is clearly established.

4 GEO-DISTRIBUTED COMPUTING FOR RESILIENCE AND CONTINUITY

4.1 Forecasting and Research Infrastructure

Numerical Weather Prediction (NWP) services at CSCS can be broadly categorized into two distinct domains: production (operational) and Research&Development (R&D), each serving unique purposes with specific requirements.

The production environment is designed to deliver weather forecast updates on a fixed schedule. This requires strict adherence to predefined timelines to ensure that critical information reaches users on time. As a result, the production setup relies on a robust and dedicated computational infrastructure that guarantees a consistent time-to-solution, the interval between data input and forecast output. Running any unrelated computational tasks on this

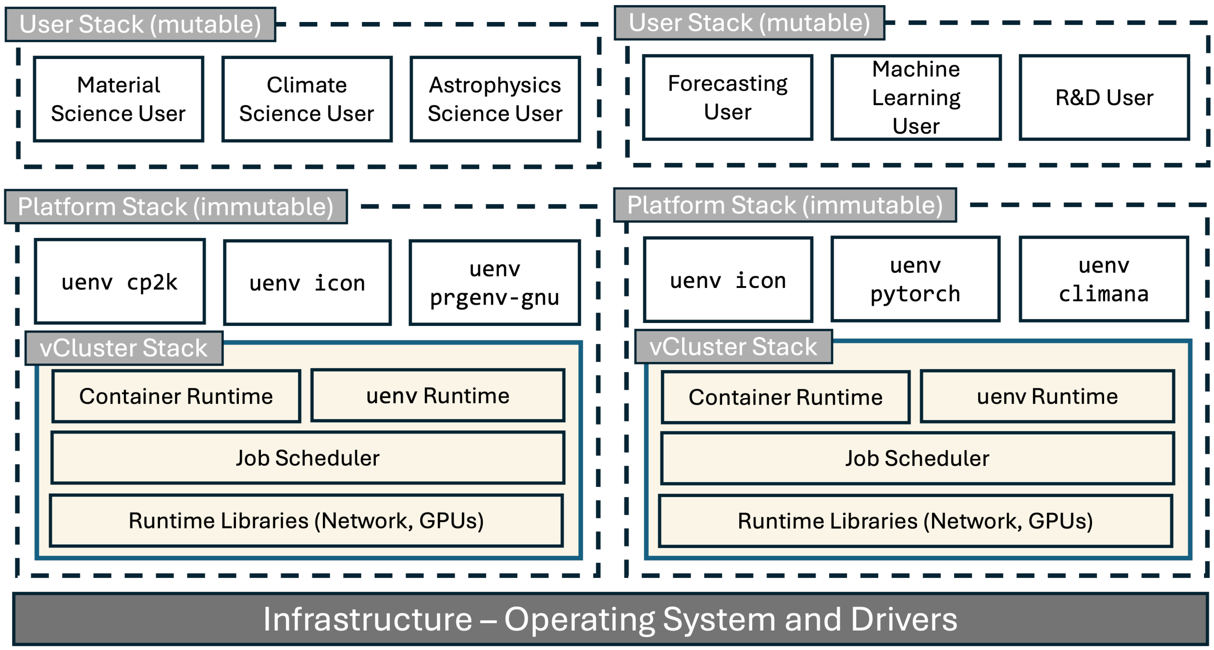

Figure 4: Depiction of two platforms on the Alps infrastructure. In this example, each platform has its dedicated platform stack, which mainly differs for the uenvs made available.

infrastructure could delay the process, potentially compromising the reliability and timeliness of forecasts, which remain essential for public safety and decision-making.

In contrast, the R&D environment is geared toward experimentation and innovation. It operates without the strict time constraints imposed on production, allowing researchers the flexibility to explore new methodologies, test software enhancements, and develop improvements in an environment that mirrors production. To prevent any impact on operational stability, R&D employs separate hardware resources - including network, storage, and compute - which also makes it possible to test applications across varied hardware platforms and supports scale resources dynamically according to project needs.

In principle, production and R&D systems could be served by entirely different configurations in both software and hardware. However, aligning the R&D setup closely with production allows successful innovations to be seamlessly transitioned into operational use without disrupting service reliability. This strategic alignment ensures that, while the production environment remains dedicated to delivering consistent and predictable results, the R&D configuration promotes the rapid adoption of novel solutions.

Historically, maintaining stability and compatibility between these environments required symmetrical, dedicated hardware for both production and R&D. Although this approach guaranteed performance and redundancy, it often led to under-utilization of R&D resources during idle periods or limited capacity during peak demand, thereby reducing overall efficiency. The evolution toward a software-defined architecture has addressed these challenges by enhancing flexibility and optimizing resource utilization. Under this model, the production environment retains fixed, dedicated resources and a stable software configuration to meet stringent reliability requirements, while the R&D environment benefits from elasticity in the resource allocation and adaptable software setups.

Moreover, the R&D environment serves as a critical fail-over mechanism for production services. During planned maintenance or unexpected system failures in the production environment, the R&D infrastructure can be quickly reconfigured into a production-ready setup to maintain service continuity and minimize downtime. This dual-purpose design highlights the importance of the infrastructure split in balancing operational reliability with innovative flexibility.

4.2 Role of the AlpsE Facility at EPFL

Thanks to the Alps infrastructure, MeteoSwiss services are now supported by two facilities: one located at the Ecole Polytechnique Federale de Lausanne (EPFL), which for clarity we refer to as AlpsE, and the main deployment of Alps in Lugano. Together, they form a geographically diverse computing environment. Rather than distributing operational and R&D services concurrently across both sites, these facilities operate in an alternating configuration. At any given time, one site is fully dedicated to operational (production) services, while the other focuses exclusively on R&D. This exclusive allocation ensures a clear separation of responsibilities, optimizes resource usage, and minimizes interference between the stringent demands of forecasting and the flexible nature of research.

The AlpsE facility plays a central role in this dynamic setup. Depending on the current configuration, AlpsE can serve either as the operational hub or as the R&D center. When designated as the operational site, AlpsE leverages its dedicated resources to meet the

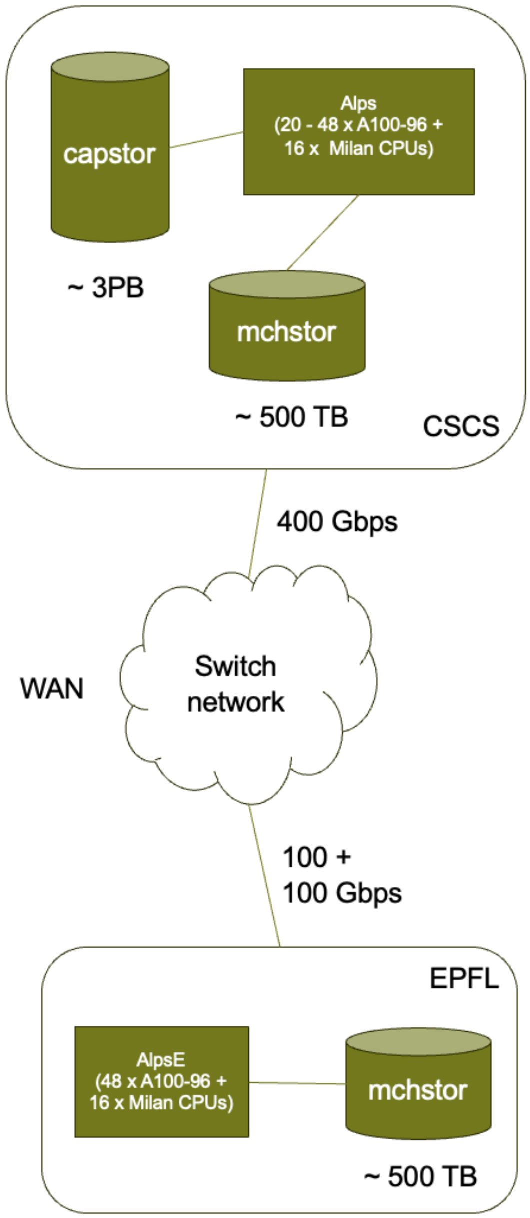

Figure 5: Resources distribution and connection between CSCS in Lugano and Alps component at EPFL.

exacting requirements of Numerical Weather Prediction (NWP) services, consistently delivering timely and reliable weather forecasts. Its infrastructure - including hardware, configuration management software, and application environments - closely mirrors that of Alps in Lugano, ensuring compatibility and a seamless transition when roles switch.

This alternating model significantly enhances resilience and adaptability. For example, during periods of planned maintenance or unexpected disruptions at the operational site, the other facility can swiftly assume production duties. The underlying software-defined architecture supports this flexibility by enabling rapid reconfiguration of the R&D site to replicate the operational setup, thereby ensuring service continuity with minimal downtime. AlpsE's strategic location at EPFL further reinforces its critical role in the fail-over strategy, whether it is actively hosting operational services or standing by to assume them from Alps in Lugano.

Importantly, access to production or R&D services is independent of their physical location. Users are abstracted from the underlying site configuration because the immutability of the software stack ensures that identical deployment images can be applied anywhere - as long as the hardware and basic operating system remain consistent. In the traditional approach (see Figure 3), maintaining identical systems across sites required explicit and often labor-intensive actions by system maintainers. With the deployment of immutable images, however, this challenge has effectively been eliminated.

4.3 High-Speed Networking and Future Geo-Redundancy

The geo-distributed configuration of the Alps infrastructure in Lausanne and Lugano presents both significant opportunities and distinct challenges, particularly in ensuring reliable fail-over and continuous service delivery. A critical requirement is the replication of data between the two facilities, which is achieved via high-speed, encrypted links. These links employ Virtual eXtensible LAN (VXLAN), a technique that encapsulates OSI layer 2 Ethernet frames within UDP datagrams [9]. Encryption is provided by a CloudSEC solution [13] to guarantee data integrity while minimizing transmission delays, thereby addressing the latency and consistency challenges inherent in a distributed setup. This high-speed networking infrastructure is essential for synchronizing operational data, allowing either site to assume production duties with minimal impact during disruptions. The SWITCH network, which supplies connectivity to research and education institutions across Switzerland, offers Wide Area Network (WAN) connectivity with 400Gbps (CSCS) and 100Gbps (EPFL) fail-over links between the two sites (see Fig. 5).

While this distributed model marks a significant advancement toward flexible and resilient service provision, achieving full geo-redundancy remains a work in progress. Beyond Numerical Weather Prediction (NWP) applications, critical base services - such as DNS and LDAP - must also be redundantly replicated across both facilities to ensure operational continuity in the event of a complete site failure. This redundancy introduces the need to carefully manage split-brain scenarios: isolated systems might operate independently, potentially creating conflicting data, while robust replication protocols are necessary to maintain consistency and prevent data conflicts. Once these foundational services are reliably in place, applications can operate seamlessly across the two sites in Lausanne and Lugano, and additional cloud providers can be integrated to further enhance scalability and resilience.

5 RESULTS

5.1 Hardware Configuration and Performance

MeteoSwiss' new infrastructure is organized into two identical vClusters: one located on the Alps system in Lugano and the other in

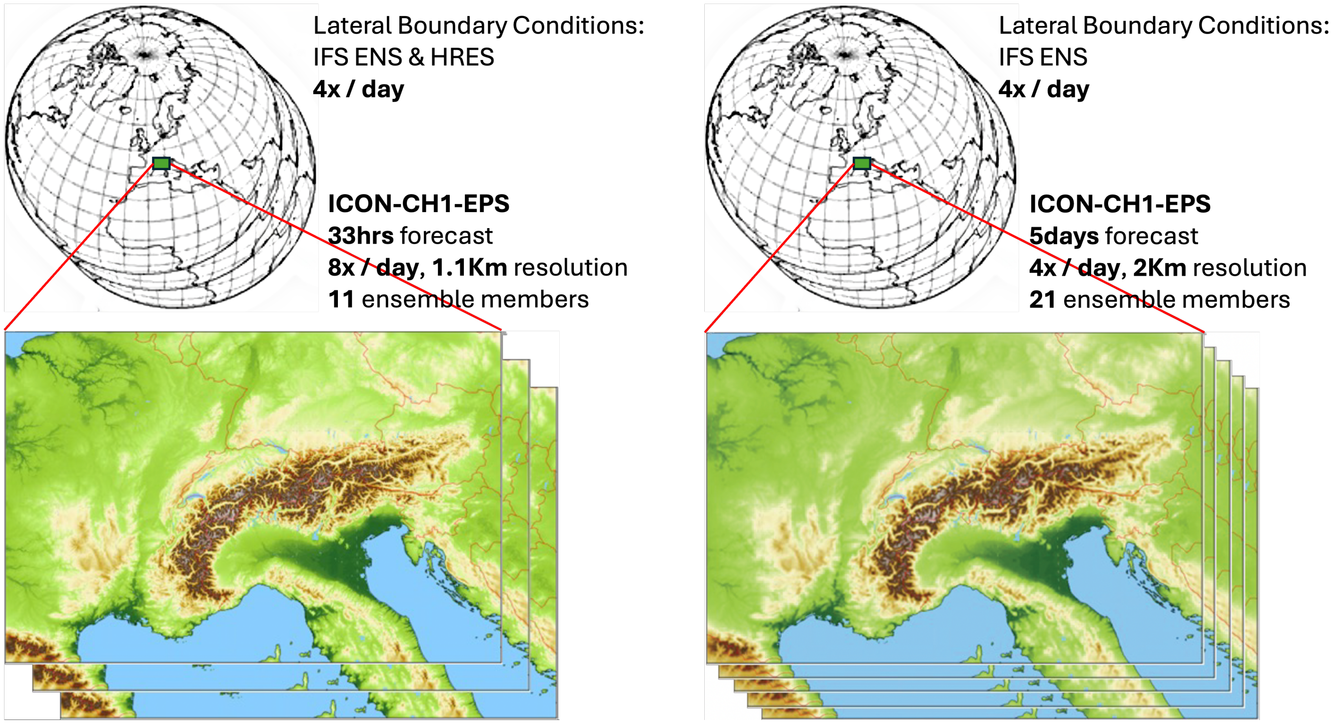

Figure 6: Set up of ICON simulations on the Alps infrastructure. Initial and lateral boundary conditions are obtained from data assimilation and larger scale and global runs. 33-hours forecasts are executed every three hours, at 1.1km resolution, with 11 ensemble members (left), while 5-days forecasts runs every 6 hours, at 2km resolution, on 21 ensemble members (right).

Lausanne. These vClusters provide compute nodes based on AMD EPYC (Zen2, MILAN) processors, each paired with four NVIDIA A100 GPUs, and are used primarily for running forecast ensembles. The compute nodes are connected via a high-speed network (HSN) using HPE Slingshot technology, offering bandwidths of up to 200 Gbps. Currently, the HSN is configured as two separate domains - one at EPFL in Lausanne and the other at CSCS in Lugano.

In addition to GPU-enabled nodes, MeteoSwiss also uses pure multicore nodes based on the same HPE architecture and AMD processor family. These nodes support ancillary services such as pre- and post-processing, as well as ensemble data assimilation - that is, the integration of sensor data, initial conditions, and lateral boundary conditions into forecast models.

Each vCluster includes a local Lustre file system that provides user home directories and, critically, scratch space for operational services. These local file systems are kept synchronized to ensure fail-over capability. In addition, a shared Lustre file system is mounted across both sites to support R&D workloads, provide long-term storage, and serve as an access point to tape archives. Both operational and R&D services use uenvs deployed consistently across the vClusters.

Beyond its geo-distributed architecture, which enhances the overall resilience of MeteoSwiss' operational services, the CSCS vCluster in Lugano is further protected by being located within the Alps system partition that is connected to an Uninterruptible Power Supply (UPS), ensuring continued service availability in case of power disruptions.

5.2 Simulation Capabilities

The operational forecasting workflow of MeteoSwiss is composed of multiple jobs, each running at different times and on different hardware partitions of the allocated infrastructure. These jobs operate under strict time-to-solution constraints, as the forecast products must be delivered under stict time constraints to downstream users.

The most computationally demanding component of the workflow is the execution of the ICON model at a resolution of 1.1 km that covers a domain covering Switzerland (see Figure6). Each forecast is executed in slightly varied configurations (ensemble runs) to provide statistically robust output.

In regular operations, two main simulation types are performed:

- **ICON-CH1-EPS:** A 33-hour forecast using 11 ensemble members, executed every three hours. Each run must complete within 55 wall-clock minutes.
- **ICON-CH2-EPS:** A 5-day forecast using 21 ensemble members, executed every six hours. Each run must complete within 45 wall-clock minutes.

Each simulation requires 8 GPUs per ensemble member. For ICON-CH2-EPS, this totals 168 GPUs (42 computing nodes) per forecast cycle.

Between ensemble forecast runs, additional tasks are executed, most notably data assimilation. These tasks incorporate real-time sensory data and lateral boundary conditions from larger-scale models - specifically the IFS model - run by European partner institutions.

5.3 Impact on MeteoSwiss Workflows

The transition from a dedicated HPC environment to the Alps infrastructure, while architecturally significant, was designed to minimize disruption to existing operational workflows. Previously, MeteoSwiss operated two software-defined clusters on isolated, dedicated infrastructure. These clusters were manually provisioned and statically configured, limiting elasticity and requiring symmetric hardware replication for R&D and operational resilience.

With the move to Alps, MeteoSwiss now operates two virtual clusters (vClusters) defined using Infrastructure-as-Code (IaC) principles and deployed on the shared, scalable Alps platform. vClusters provide a controlled and isolated execution environment that replicates the operational semantics of the previous dedicated clusters. As a result, the transition has preserved the user experience and workflow interface, avoiding the need for substantial retraining or re-engineering of forecasting pipelines. Despite this continuity, the new architecture introduces several functional improvements that impact workflows:

- Isolated execution on a shared infrastructure: vClusters provide strict namespace and resource separation, enabling MeteoSwiss to operate predictably within a multi-tenant environment without interference from other platforms.
- Improved test and validation capability: The use of `uenvs` has streamlined environment management and facilitated reproducibility. These environments can be versioned and deployed alongside production systems, simplifying the validation of model updates or system changes. In the future, support for deploying test vClusters will further enhance staging and pre-production testing workflows.
- Dynamic resource allocation: On-demand access to resources across the broader Alps infrastructure allows MeteoSwiss to scale compute resources beyond the fixed capacity of the previous system. This elasticity is especially beneficial during forecast-critical events or during parallel development and validation phases.
- Simplified fail-over and resiliency: The alignment of production and R&D vClusters across geographically distinct facilities (CSCS and EPFL) enables rapid role switching during planned maintenance or failure scenarios, improving service continuity with minimal configuration changes.

These changes improve the flexibility, scalability, and resilience of MeteoSwiss's operational workflows without disrupting their core structure. The transition demonstrates that vCluster and `uenv` technologies can preserve established practices while incrementally enabling more advanced deployment, testing, and fail-over strategies

6 CONCLUSION AND FUTURE WORK

The evolution of MeteoSwiss's Numerical Weather Prediction (NWP) system marks a significant advancement toward a more adaptive, resilient, and future-ready forecasting infrastructure. The transition to the ICON model, combined with the adoption of the Alps platform, has modernized operational forecasting by introducing scalable, flexible computing resources and laying a robust foundation for future developments in meteorology and climate science.

The underlying software-defined infrastructure enables a clear separation of compute and data services, reducing cross-dependencies and simplifying system maintenance. This flexibility is especially critical as emerging technologies - such as AI - become increasingly important in weather prediction. The current architecture positions MeteoSwiss to stay at the forefront of innovation in the field.

Looking ahead, MeteoSwiss will continue to track and adopt developments in vCluster and `uenv` technologies, which are designed to enhance automation and clarify the delegation of operational responsibilities. These innovations will play a key role in transforming the existing geo-distributed infrastructure into a geo-redundant system - one capable of automating a growing number of service continuity functions.

MeteoSwiss will continue to follow the evolution of vCluster and `uenv` technologies, which aim to increase automation and enable clearer delegation of responsibilities. These advancements will be instrumental in evolving the current geo-distributed solution into a geo-redundant forecasting system, capable of automating an increasing number of service continuity aspects.

ACKNOWLEDGMENTS

We dedicate this paper to André Walser, whose unwavering support and deep insight were instrumental in making the implementation described here possible. Though he has passed away, the foundation he helped establish will continue to shape future developments for many years to come.

The authors thank Maxime Martinasso and Ben Cumming (CSCS) for their valuable input and discussions, which helped shape this manuscript.

An AI-generated tool based on ChatGPT built on GPT-4 architecture has been used to enhance the readability of all sections of this document. The authors have carefully integrated the AI-suggested edits to preserve the intended meaning. The AI tool was not used to generate ideas or data.

REFERENCES

[1] Sadaf R Alam, Miguel Gila, Mark Klein, Maxime Martinasso, and Thomas C Schulthess. 2023. Versatile software-defined HPC and cloud clusters on Alps supercomputer for diverse workflows. *The International Journal of High Performance Computing Applications*, 37, 3-4, 288–305. DOI: 10.1177/10943420231167811.

[2] Sadaf R. Alam, Mark Klein, Mauro Bianco, Roberto Aielli, Xavier Lapillonne, Andre Walser, and Thomas C. Schulthess. 2020. Software defined infrastructure for operational numerical weather prediction. In *Driving Scientific and Engineering Discoveries Through the Convergence of HPC, Big Data and AI*. Jeffrey Nichols, Becky Verastegui, Arthur 'Barney' Maccabe, Oscar Hernandez, Suzanne Parete-Koon, and Theresa Ahearn, (Eds.) Springer International Publishing, Cham, 303–317. ISBN: 978-3-030-63393-6.

[3] M. Baldauf, A. Seifert, J. Förstner, D. Majewski, M. Raschendorfer, and T. Reinhardt. 2011. Operational convective-scale numerical weather prediction with the COSMO model: description and sensitivities. *Mon. Wea. Rev.*, 139, 3887–3905.

[4] Luca Bonaventura and Todd Ringler. 2005. Analysis of discrete shallow-water models on geodesic delaunay grids with c-type staggering. *Monthly Weather Review*, 133, 8, 2351–2373. DOI: 10.1175/MWR2986.1.

[5] Jonathan Coles, Ben Cumming, Theofilos-Ioannis Manitaras, Jean-Guillaume Piccinali, Simon Pintarelli, and Harmen Stoppels. 2023. Deploying alternative user environments on Alps. In *CUG '23*. https://api.semanticscholar.org/CorpusID:274643112.

[6] Felipe Cruz, Manuel Sopena Ballesteros, and Alejandro J. Dabin. 2024. Deploying Cloud-Native HPC Clusters on HPE Cray EX. In *CUG '24'*. https://api.semanticscholar.org/CorpusID:274658355.

[7] O. Fuhrer, C. Osuna, X. Lapillonne, T. Gysi, B. Cumming, M. Bianco, A. Arteaga, and T. C. Schulthess. 2014. Towards a performance portable, architecture agnostic implementation strategy for weather and climate models. *Supercomput. Front. Innovations*, 1, 45–62.

[8] X. Lapillonne and O. Fuhrer. 2014. Using compiler directives to port large scientific applications to GPUs: an example from atmospheric science. *Parallel Process. Lett.*, 24.

[9] Mallik Mahalingam, Dinesh Dutt, Kenneth Duda, Puneet Agarwal, Larry Kreeger, T. Sridhar, Mike Bursell, and Chris Wright. 2014. Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks. RFC 7348. (Aug. 2014). DOI: 10.17487/RFC7348.

[10] Riccardo Di Maria et al. 2025. Infrastructure as a service with strong tenant separation on a supercomputer. In *Proceedings of the Cray User Group Conference (CUG 2025)*. To be published. Cray User Group. New Jersey, USA, (May 2025).

[11] Maxime Martinasso, Mark Klein, Benjamin Cumming, Miguel Gila, Felipe Cruz, Alberto Madonna, Manuel Sopena Ballesteros, Sadaf R. Alam, and Thomas C. Schulthess. 2024. Versatile software-defined cluster for HPC using cloud abstractions. *Computing in Science & Engineering*, 26, 3, 20–29. DOI: 10.1109/MCSE.2024.3394164.

[12] Maxime Martinasso, Mark Klein, and Thomas C. Schultess. 2025. Alps, a versatile research infrastructure. In *Proceedings of the Cray User Group Conference (CUG 2025)*. To be published. Cray User Group. New Jersey, USA, (May 2025).

[13] Avijit Mondal and Pinaki Chatterjee. 2024. CloudSec: a lightweight and agile approach to secure medical image transmission in the cloud computing environment. *SN Computer Science*, 5, (Jan. 2024). DOI: 10.1007/s42979-023-02539-w.

[14] S. Schuppli et al. 2025. Evolving HPC services to enable ML workloads on HPE Cray EX. In *Proceedings of the Cray User Group Conference (CUG 2025)*. To be published. Cray User Group. New Jersey, USA, (May 2025).

[15] J. Steppeler, G. Doms, U. Schättler, H. Bitzer, A. Gassmann, U. Damrath, and G. Gregoric. 2003. Meso-gamma scale forecasts using the nonhydrostatic model LM. *Meteor. Atmos. Phys.*, 82, 75–96.

[16] EXCLAIM Project Team. 2025. EXCLAIM: Exascale Computing Platform for Cloud-Resolving Weather and Climate Models. https://exclaim.ethz.ch/. Accessed: 2025-04-05. (2025).

[17] Andy B. Yoo, Morris A. Jette, and Mark Grondona. 2003. SLURM: simple linux utility for resource management. In *Job Scheduling Strategies for Parallel Processing*. Dror Feitelson, Larry Rudolph, and Uwe Schwiegelshohn, (Eds.) Springer Berlin Heidelberg, Berlin, Heidelberg, 44–60. ISBN: 978-3-540-39727-4.

[18] G. Zängl, D. Reinert, P. Rípodas, and M. Baldauf. 2015. The ICON (ICOsahedral non-hydrostatic) modelling framework of DWD and MPI-M: description of the non-hydrostatic dynamical core. *Quarterly Journal of the Royal Meteorological Society*, 141, 563–579.