Connecting ECMWF Workflows to the Future: Leveraging Distributed Pre-Exascale Resources

Johannes Bulin*, Cristina Duma†, Matthew Griffith§, Ulrike Hager*, Bentorey Hernandez-Cruz*, Gareth Jones*, Emma Kuwertz§, Julien Lefaucheur§, Giuseppe Misurelli†

*Research Department/Integrated Forecasting System, †Computing Department/Application Delivery, §Forecasts & Services Department/Production Services

Destination Earth

In Destination Earth we are using EuroHPC supercomputers to run high resolution Digital Twins (DTs).

What if…

- our workflows could **seamlessly scale** across **distributed HPC systems**?
- users could tap into pre-exascale systems through **familiar intuitive tools**?
- we could **monitor and understand** every component, **no matter where it runs**?

The Destination Earth Workstreams

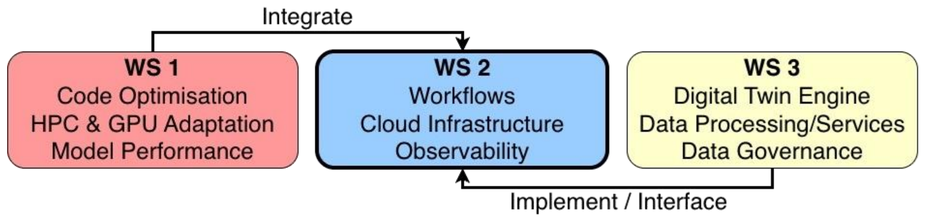

In WS2 **Agile practices** keep teams **closely connected** for fast, aligned collaboration.

Extremes Digital Twin Orchestration on EuroHPC

The **Global Weather-Induced Extremes DT** is operated by ECMWF on the **LUMI Supercomputer** in Finland.

- **Daily** forecasts at **4.4 km**
- Data made available on **remote FDB** on **Data Bridges**
- ECMWF **software modules installation** on **remote HPC**
- ecFlow servers hosted on the Data Bridges for **production** and **user experiments**
- **Secure connections** to remote HPC via ssh

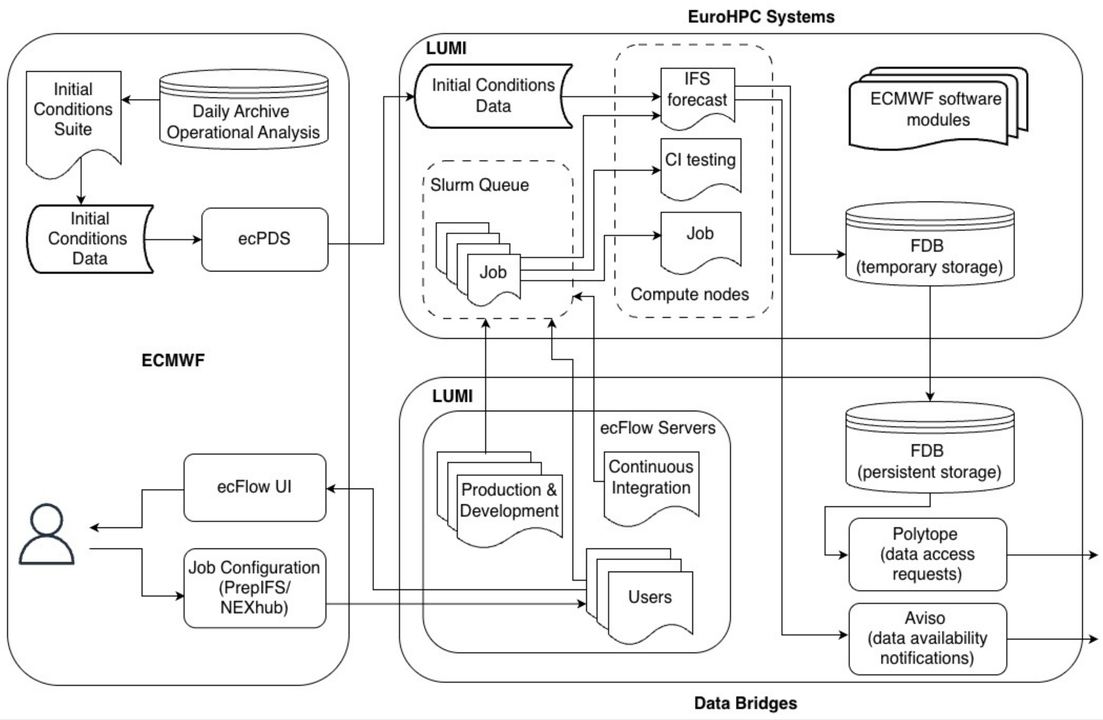

- EuroHPC porting work implemented in dedicated **IFS cycles 48R3** and **49R3**
- Robust **automated testing** pipeline on both Atos and LUMI
- Developments to be **contributed back** into core IFS cycles

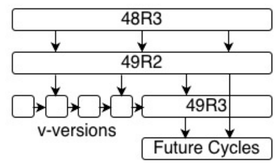

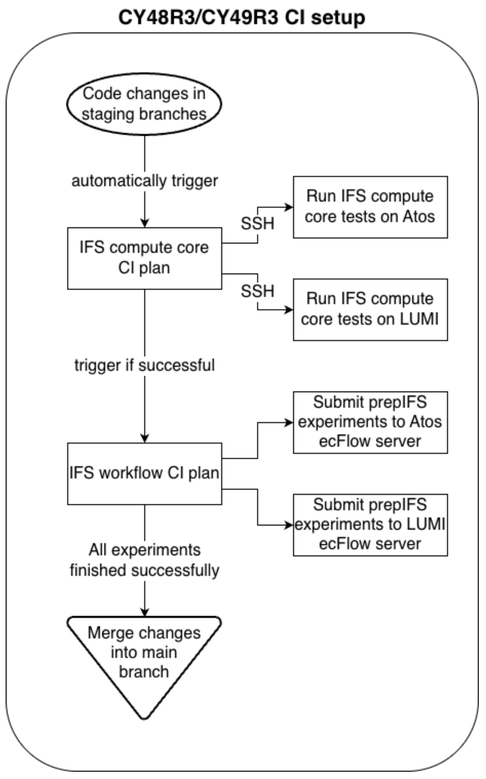

Cloud Platform Management and Observability

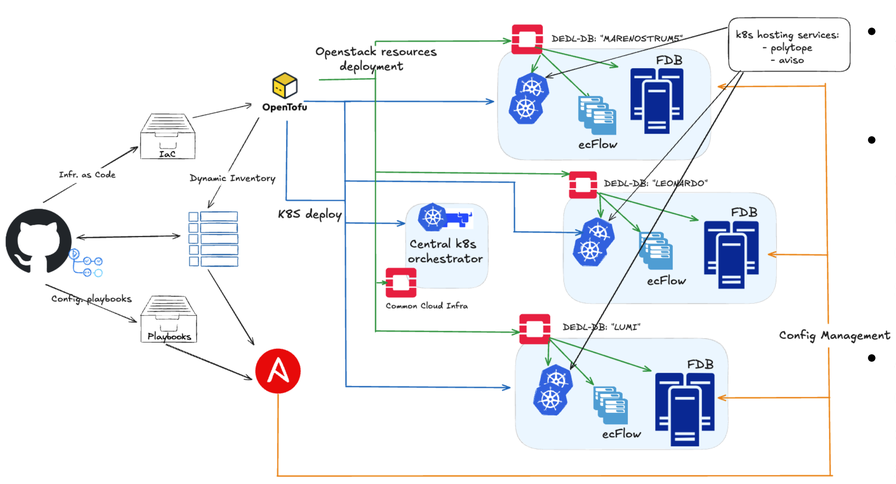

- **Automated deployment** via OpenTofu + Ansible.
- Dynamic inventory & configuration management for **consistent setups across sites**.
- A **shared Kubernetes control plane** coordinates deployments and shared services.

- Each site gathers metrics and logs via a local collector for **distributed data collection**
- Observability data **delivered to ECMWF** for unified indexing and long-term storage
- Grafana **dashboards** provide **real-time status**, **performance**, and **health across all sites**
- **Automated alerts** highlight failures, delays, or certificate issues

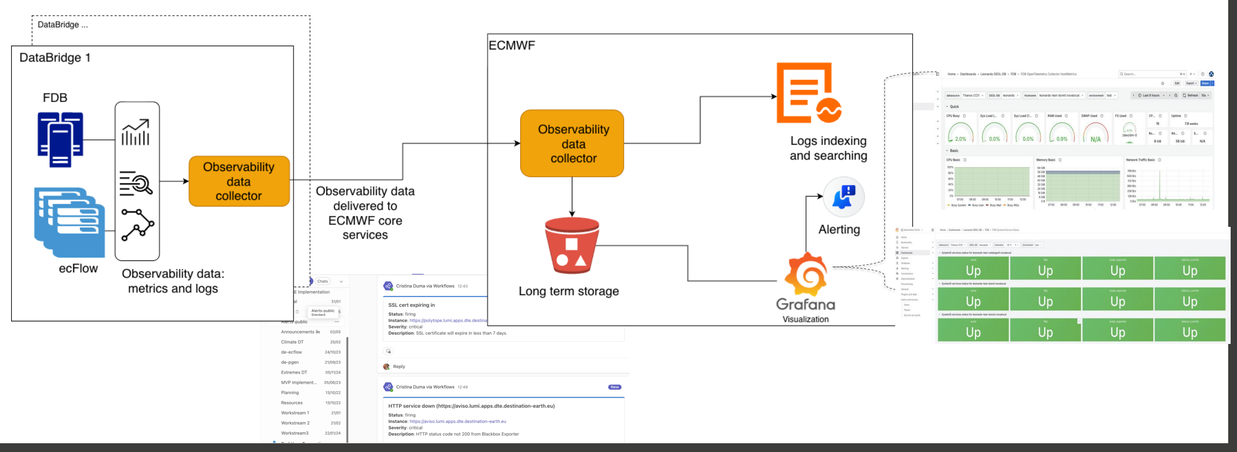

Technical challenges

Security
- **Challenge:** Security rules, certificate lifecycle management, internet access
- **Solution:** HashiCorp Vault-based keys, secure proxies for secrets/certificates

Hardware Diversity
- **Challenge:** Mixed AMD / NVIDIA / INTEL and CPU / GPU combinations
- **Solution:** Automated testing, OpenMP/OpenACC, Loki

Service deployment
- **Challenge:** Provision / configure across distributed and complex systems
- **Solution:** Infrastructure / configuration as code (Terraform / Ansible)

Software and Data
- **Challenge:** Module environment differs across sites, no MARS access
- **Solution:** Easybuild-based modules on EuroHPC, Polytope, ecPDS

Observability
- **Challenge:** Remote metrics must be captured securely
- **Solution:** Exporters → central monitoring, secure ingestion

Operations
- **Challenge:** No control over HPC operations, inconsistent resource reporting
- **Solution:** Relationship building and negotiation with sites

Data Governance
- **Challenge:** Requirements for full GRIB2 outputs pre-2024
- **Solution:** Bespoke data layout and early MultIO integration for encoding

Key Achievements

- First ever **daily production runs** on an **external system** (LUMI) – at 4.4 km
- **Cloud tenancy** to host distributed services deployed on **three Data Bridges**
- Securing ecFlow servers for **flexible suite deployment** on LUMI and LEONARDO, interfacing with prepIFS / NEXhub and prepML
- **Integrating observability metrics** into ECMWF Observability Platform

We've turned ideas into real results and laid the groundwork to take full advantage of future exascale systems.

Progress Through Collaboration

WS2 is **distributed across teams and sites**, fostering **collaborations and interactions** which would otherwise not exist.

The progress seen to date would not have been possible without the close collaboration with WS1 and WS3, the km-scale community at ECMWF and teams across the centre responsible for key software we rely upon – ecflow, FDB, IFS, MultIO, Polytope, Aviso to name a few.

Collaboration and transparency have been key to the continued success of the project.