



The Bureau
of Meteorology

Past, present and future of HPC at the Australian Bureau of Meteorology

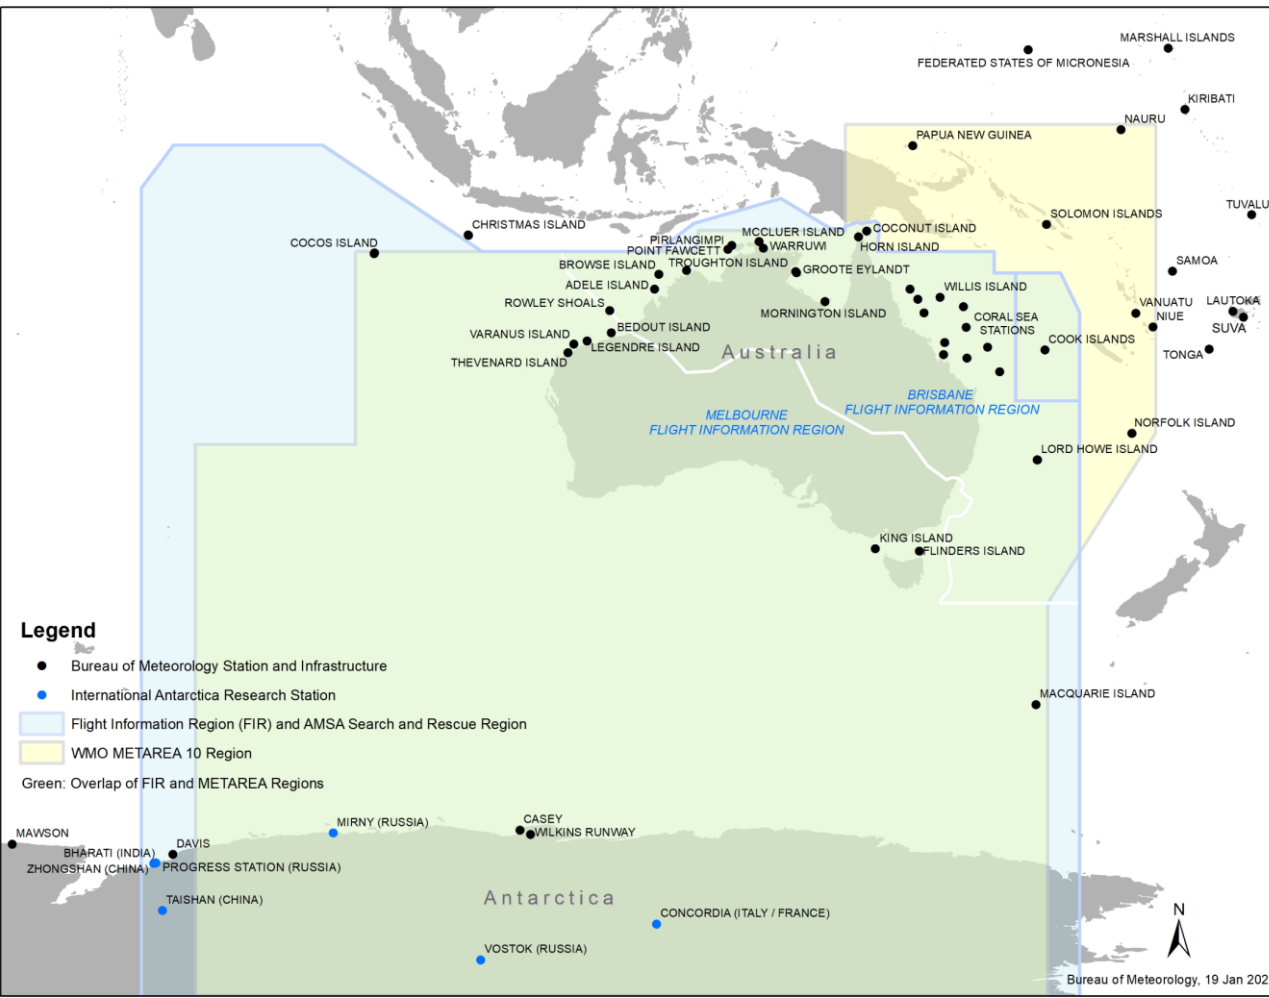
Tom Gale

21st ECMWF workshop on High Performance Computing in Meteorology

2025-09-18

Bureau of Meteorology

- Weather, water, climate, ocean and space
- Large area coverage – Australian mainland, Tasmania, Antarctica, Indian and Pacific oceans
- Many Australians get weather information and warnings direct via our website and phone apps



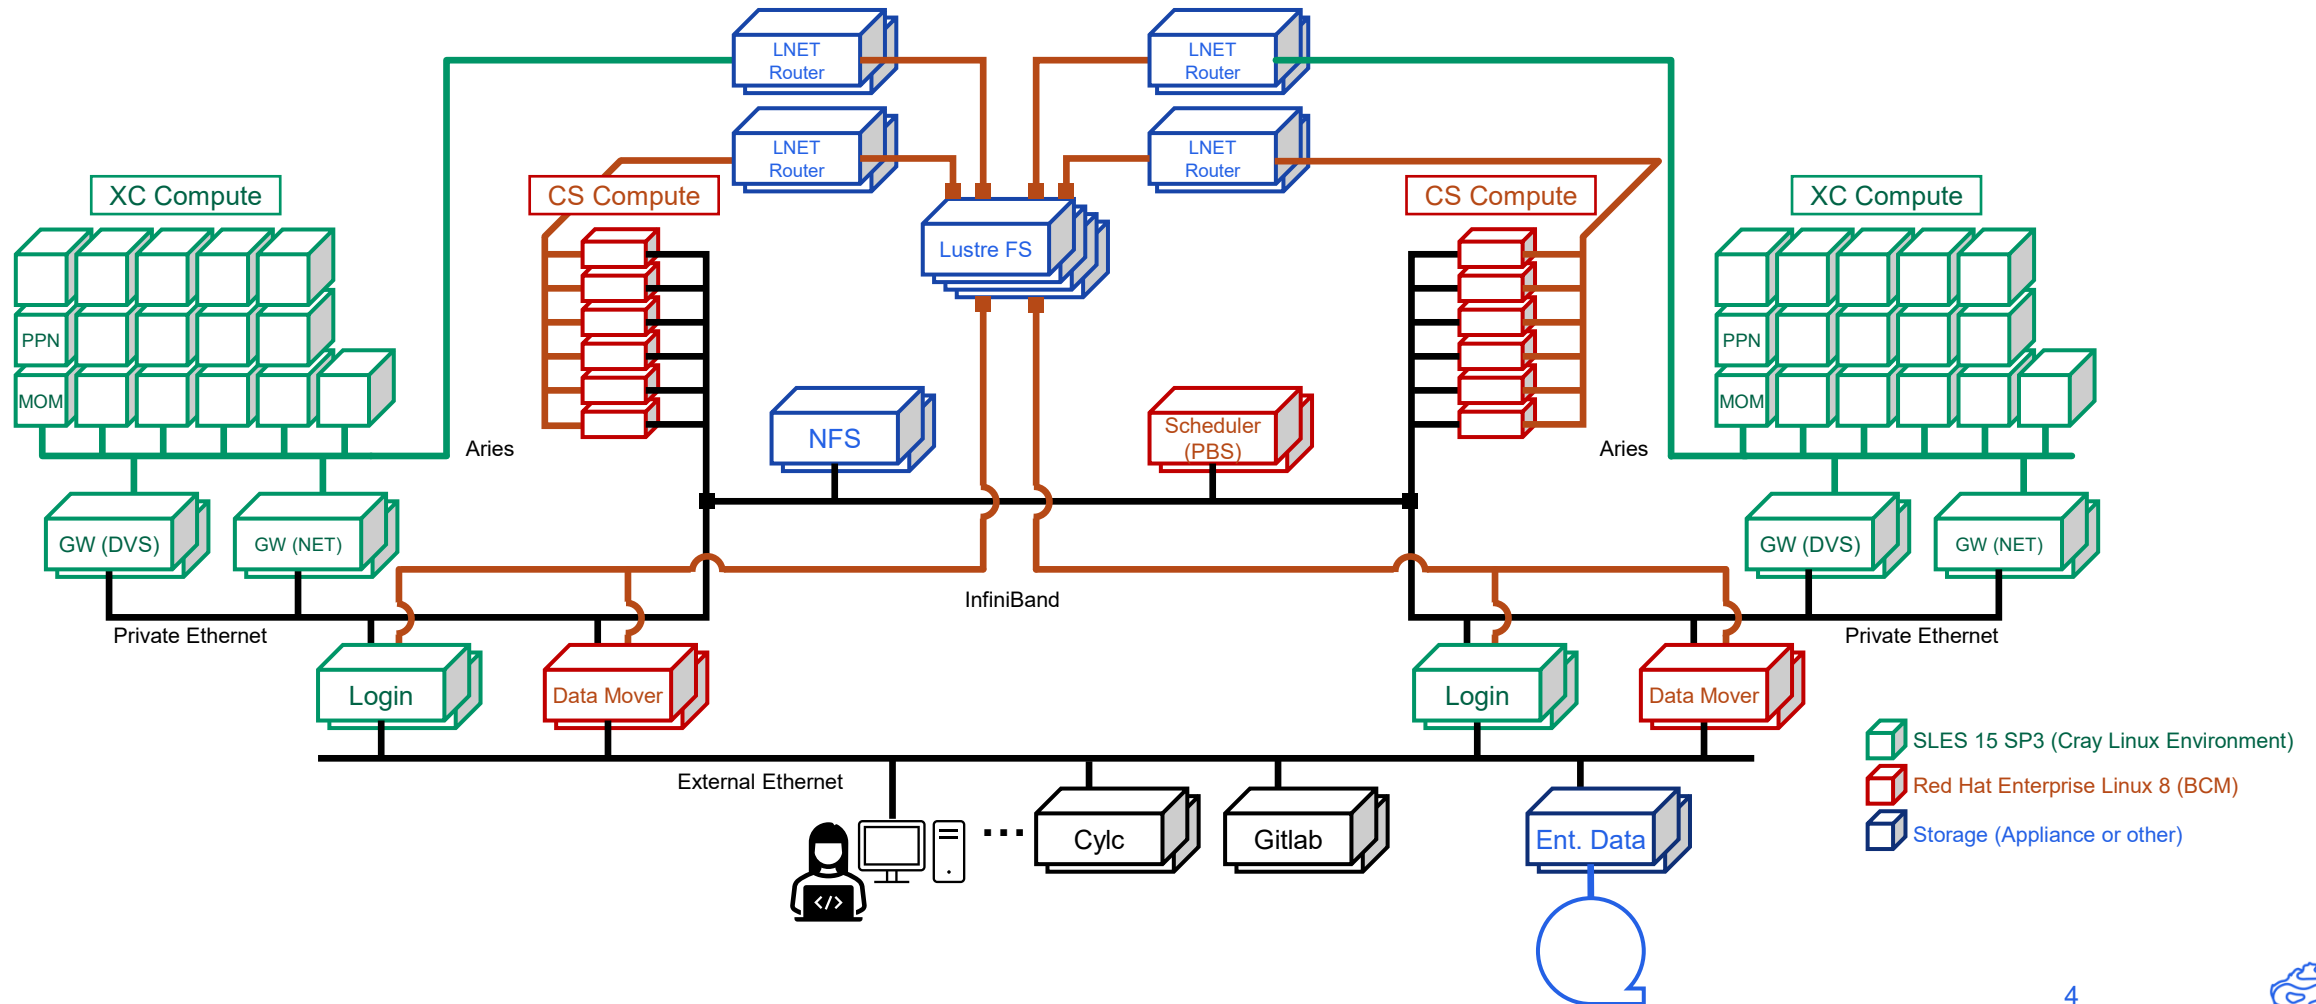
The past (since 20th ECMWF HPC workshop 2023)

- Decommissioned Australis - Cray XC40 (~2200 nodes Intel Haswell) and associated Lustre storage
 - 9 years in production
 - Consistently >99.9% per month availability for production
- Decommissioned Aurora - Cray CS400 (~40 nodes Intel Broadwell) and associated GPFS storage
- Decommissioned very old tape robot
- Exited co-located data centre



The present – Australis II

- Cray XC50 and CS500 in two-halves for high availability
- ~2000 nodes, Intel Skylake
- 4x 5 PB Lustre filesystems – two tuned for bandwidth, two tuned for smaller files



The present – Gadi

- Research, external collaboration and proof of concept work on Australia's National Computational Infrastructure (NCI) system Gadi
- Research work transitioning to operations is ported to internal Bureau of Meteorology systems
- Research data sharing with universities and partners
- ~10% share of ~5000 node HPC system
- Mix of Intel Cascade Lake, Intel Sapphire Rapids, Nvidia Volta, Nvidia Ampere



The present – Australis II surrounding infrastructure

- New datacentre “EDC” (commercial, co-located)
- New network with better security features
- Managed file transfer service
- Mobius (Gitlab CI)
- Kit (Artifactory with internal conda-forge mirror)

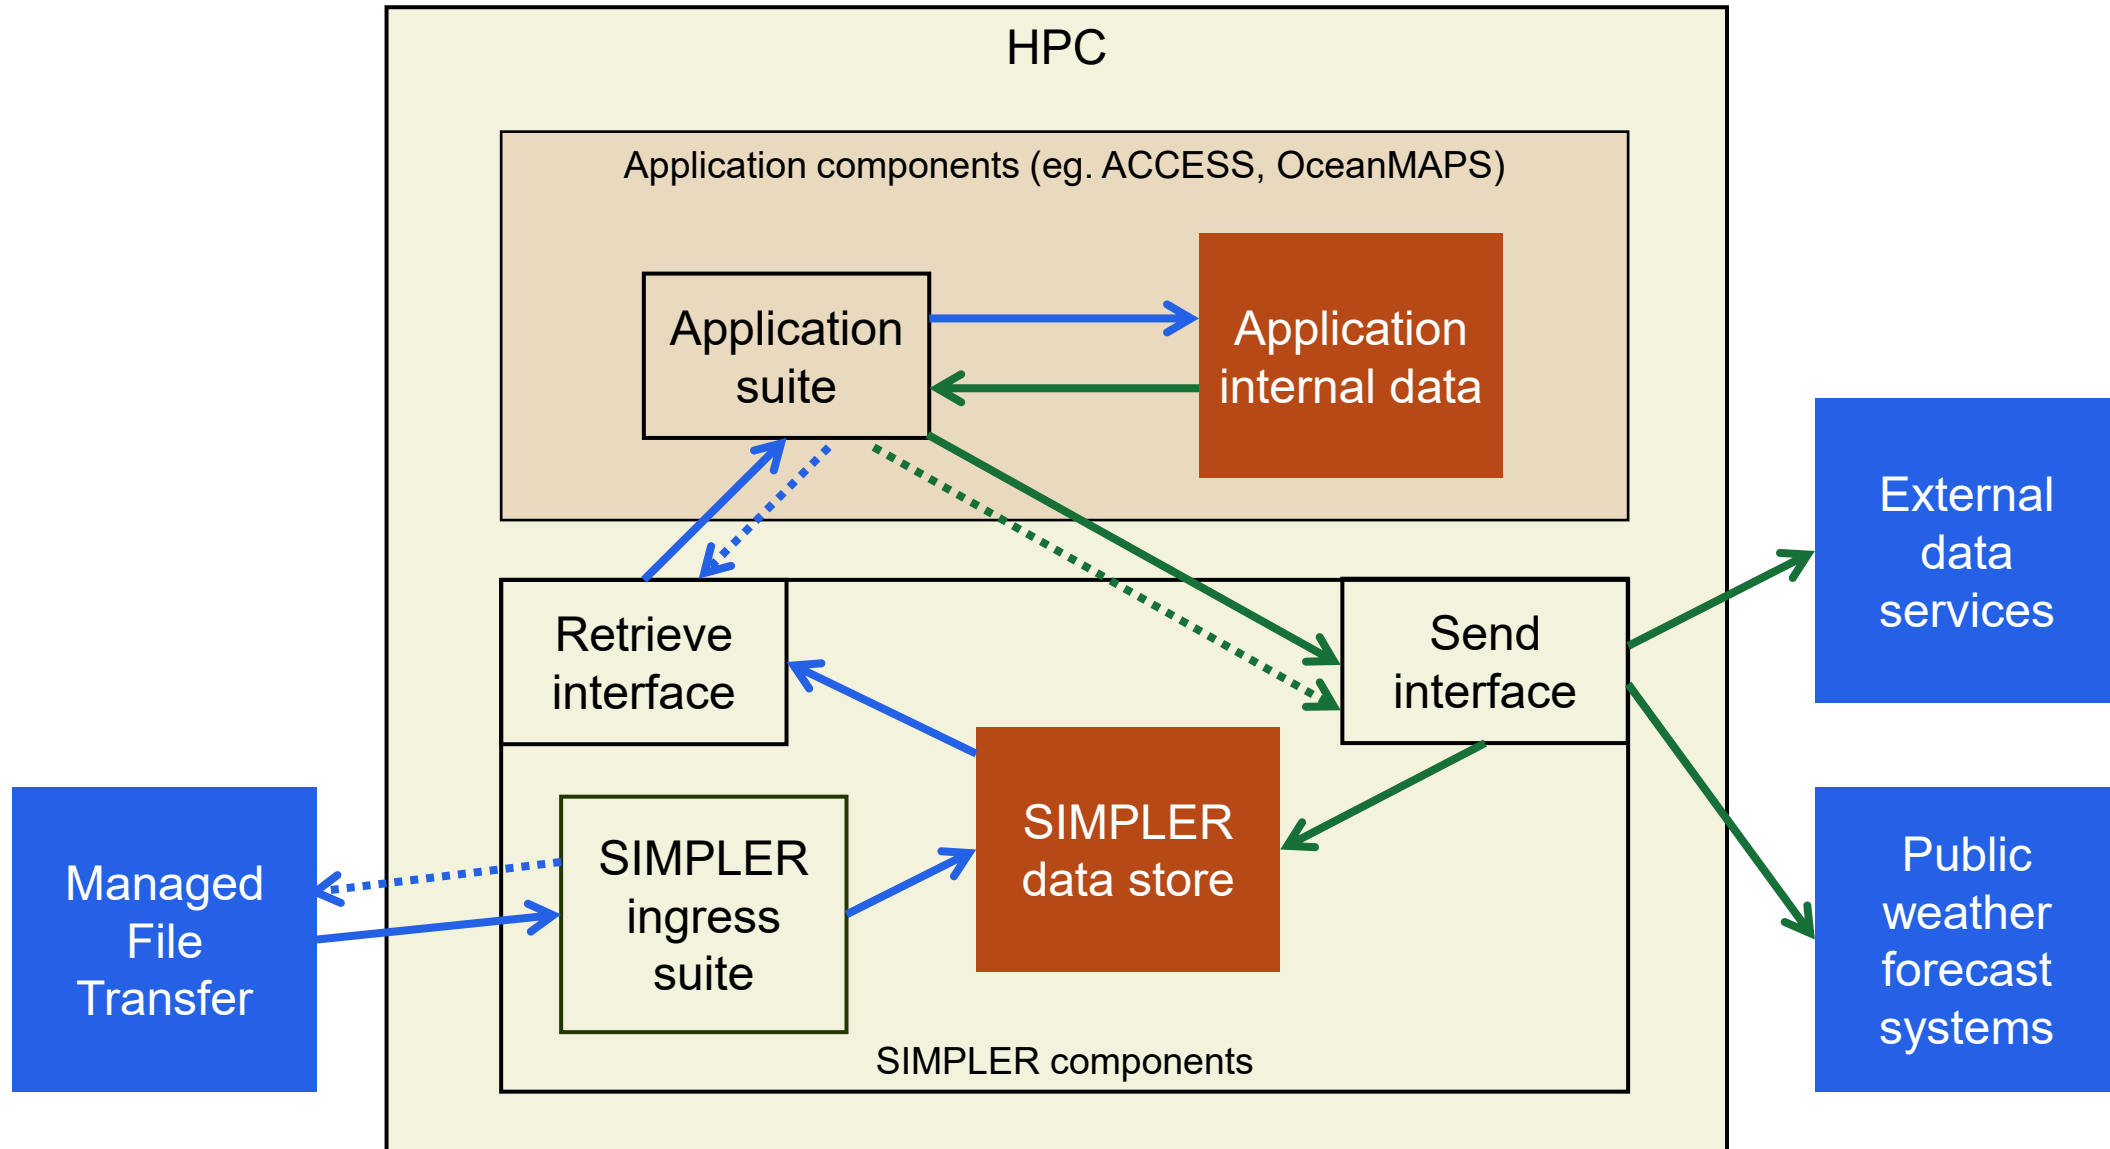


The present – models and applications

- ACCESS global, global ensemble, city (x7), city ensemble (x7) and seasonal – all Unified Model based
- OceanMAPS global ocean model
- AUSWAVE global wave model
- IMPROVER probabilistic multi-ensemble blend is the primary data source for public weather forecasts
- Dispersion models
- Storm Surge
- SIMPLER – abstract sources and destinations from models consuming and producing data
- Mixed workload on HPC
 - Big and small models
 - Pre and post processing
 - Production and non-production



SIMPLER




The near future - Sentinel

- Cray EX4000 in two halves for high availability
- ~900 nodes, Intel Sapphire Rapids, about double cores per node of Australis II XC50
- Similar core count to Australis II but +50% potential PFLOPs
- 2x 5PB Lustre filesystems
- Disaster recovery for production, otherwise available for research workload



The near future – Sentinel surrounding infrastructure

- Second new datacentre “FDC” (also commercial co-located)
 - Physically separated (interstate, hundreds of km) from EDC
 - Dedicated private fibre link between sites
- Dell PowerScale high reliability clustered storage for warm start and other recovery data, external to both HPC systems
- New Spectra Logic tape robot with migration of data from previous robot via Versity



The near future – models and applications

- National Analysis System (NAS)
 - Unified Model based hourly cycling analysis, variety of observation fittings
- ACCESS-A/AE
 - Unified Model at ~2km resolution covering Australian continent and surrounds, with ensemble
- AI/ML models
 - Currently on Gadi
- SIMPLER
 - Expand role to standardise data and metadata

The future – HPC Plan 2025-29


Six major topics:

- Investment planning
- Model enablement
- Data management
- Support services
- Disaster recovery
- Environmental sustainability



The future

- Enterprise data fabric for sharing between HPC sites and large scale data distribution to forecast systems
- Data catalogue built on Zarr datasets



The future

- Maintaining disaster recovery capability across HPC upgrade cycle
- Expect HPC systems to be more diverse
 - Visualisation and data analysis
 - Support AI/ML models in production
 - Future role of large language models in public weather forecasts and warnings?

Thank you

