Canadian Meteorological Centre – 2025 Update

WGNE40

Ron McTaggart-Cowan, Stephane Gaudreault, Syed Husain, Leo Separovic, and Christopher Subich

Environment and Climate Change Canada
Environnement et Changement climatique Canada

Overview

- Operational system upgrades in 2025
  - Deployment of a hybrid global deterministic prediction system
- Ongoing experiments with machine learning
  - Development and training of the GEML AI model
  - Exploration of AIWP with Paradis
- Future improvements
  - Computing upgrade in 2026
  - Physical model development (yes, we still do that)
  - Operational upgrade planned for 2027

An Operational Hybrid Model

- A hybrid physical-AI modelling system has been running with an “experimental” classification in Canadian Meteorological Centre operations since June 2025
- The Global Environmental Multiscale (GEM) physical model is spectrally nudged to inferences from the Global Environmental Machine Learning (GEML) model:
  - Forecasts of the large-scale state benefit from AI predictions
  - A full suite of physically consistent outputs is produced (fcst and evaluation)

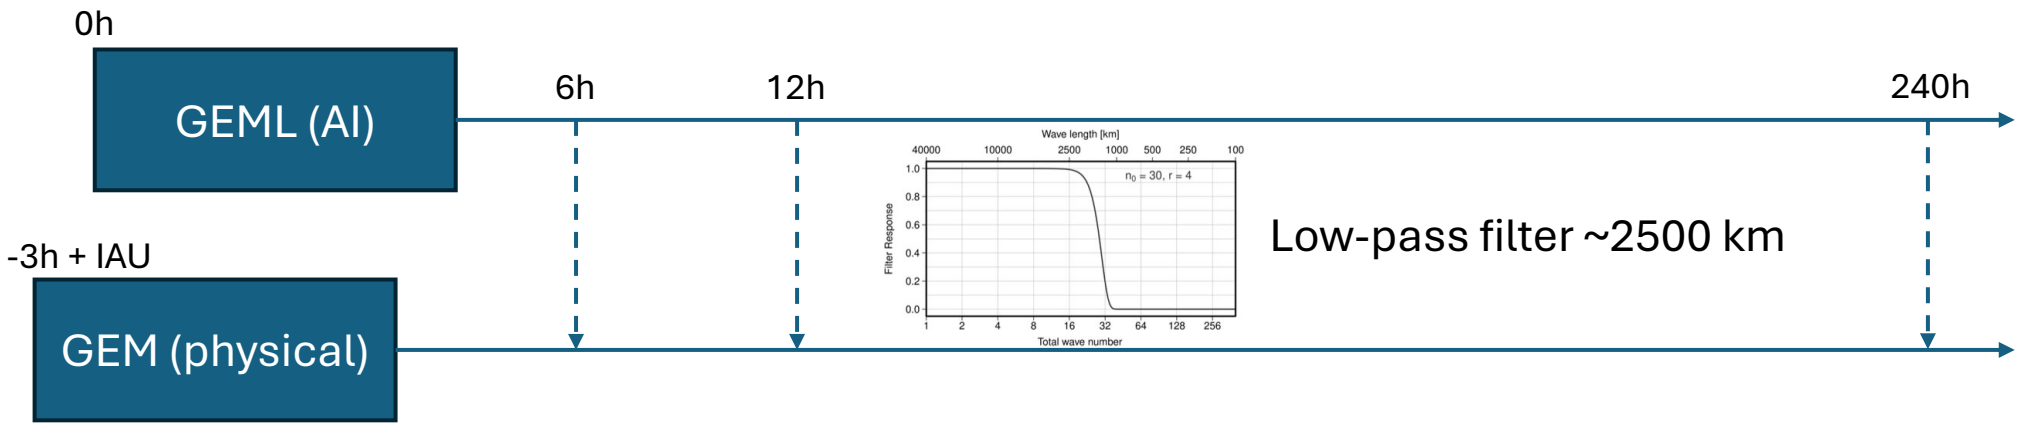

Hybrid Model Configuration

GEM (physical)
GEML (AI)
Hybrid

- The 2500 km cutoff for spectral nudging was chosen to retain scales at which GEML has a full energy spectrum
- Initial tests with smaller cutoffs showed improved RMSE; however, the fields were found to be variance-deficient at the synoptic scales
- With a Yin-Yang grid, the spectral decomposition is implemented with a pair of discrete cosine transforms
- Nudging is applied only to T, **V** in the free troposphere, with interpolation for GEML 17 pressure levels
- Time-varying nudging limits “ensemble” smoothing, with τ=3.3 equivalent to 12-h of uniform relaxation

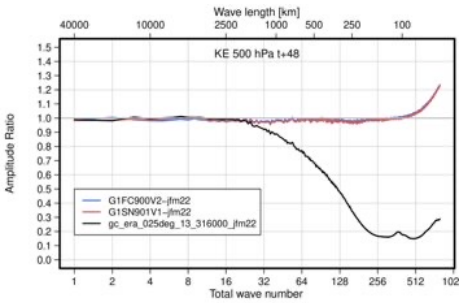

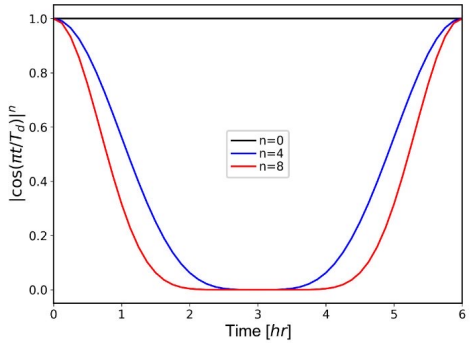

Spectral energy amplitude ratio (top) at 500 hPa after 2 days. Temporal evolution of nudging coefficients (bottom).

Hybrid Model Configu

- The 2500 km cutoff for spectral … to retain scales at which GEML ha… um
- Initial tests with smaller cutoffs … RMSE; however, the fields were f… variance-deficient at the synoptic scales
- With a Yin-Yang grid, the spectral decomposition is implemented with a pair of discrete cosine transforms
- Nudging is applied only to T, **V** in the free troposphere, with interpolation for GEML 17 pressure levels
- Time-varying nudging limits “ensemble” smoothing, with τ=3.3 equivalent to 12-h of uniform relaxation

Reduced variance in GEML at long lead times extends beyond 2500 km, reducing hybrid-model kinetic energy by ~10%.

GEM (physical)
GEML (AI)
Hybrid

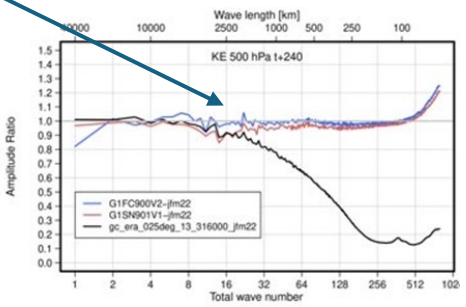

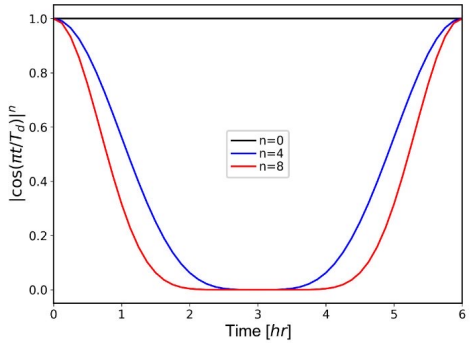

Spectral energy amplitude ratio (top) at 500 hPa after 10 days. Temporal evolution of nudging coefficients (bottom).

Hybrid Model Deployment

Summer 2022 Period (120h)

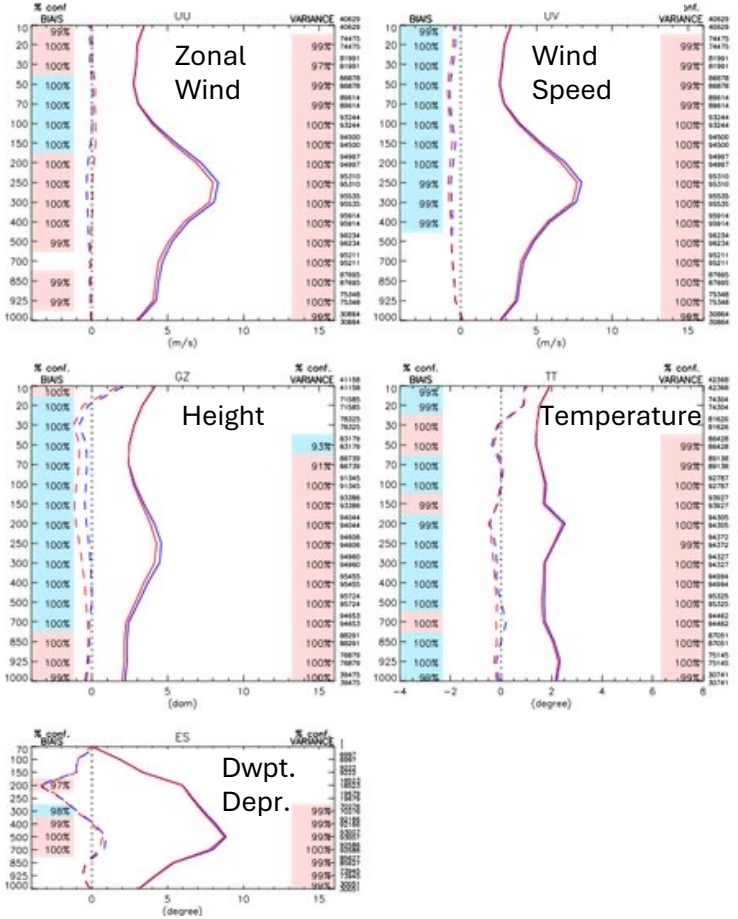

Winter 2022/23 Period (120h)

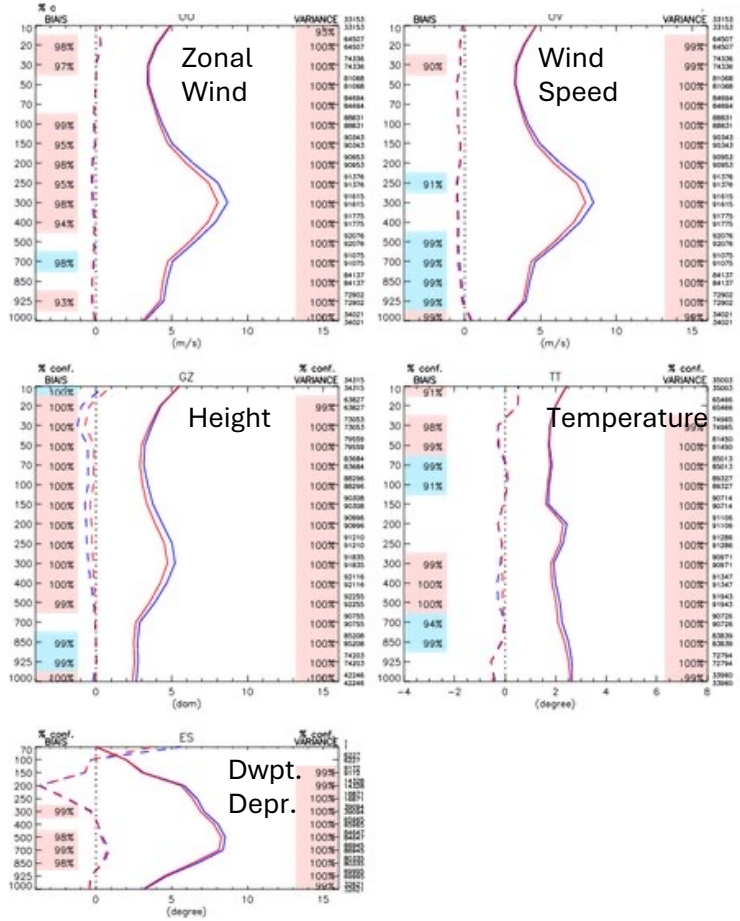

Forecast improvement is equivalent to a major physical-model upgrade (5+ years of research).

Spectral nudging to GEML is **not** simply bias correction: flow-dependent error is significantly reduced.

Improvements extend to non-nudged variables and levels.

Bias (dashed) and error std dev (solid) of the operational GDPS (blue) and hybrid model (red) against global radiosondes for 2.5-month periods. Significant changes are identified with boxes.

Tropical Cyclone Prediction

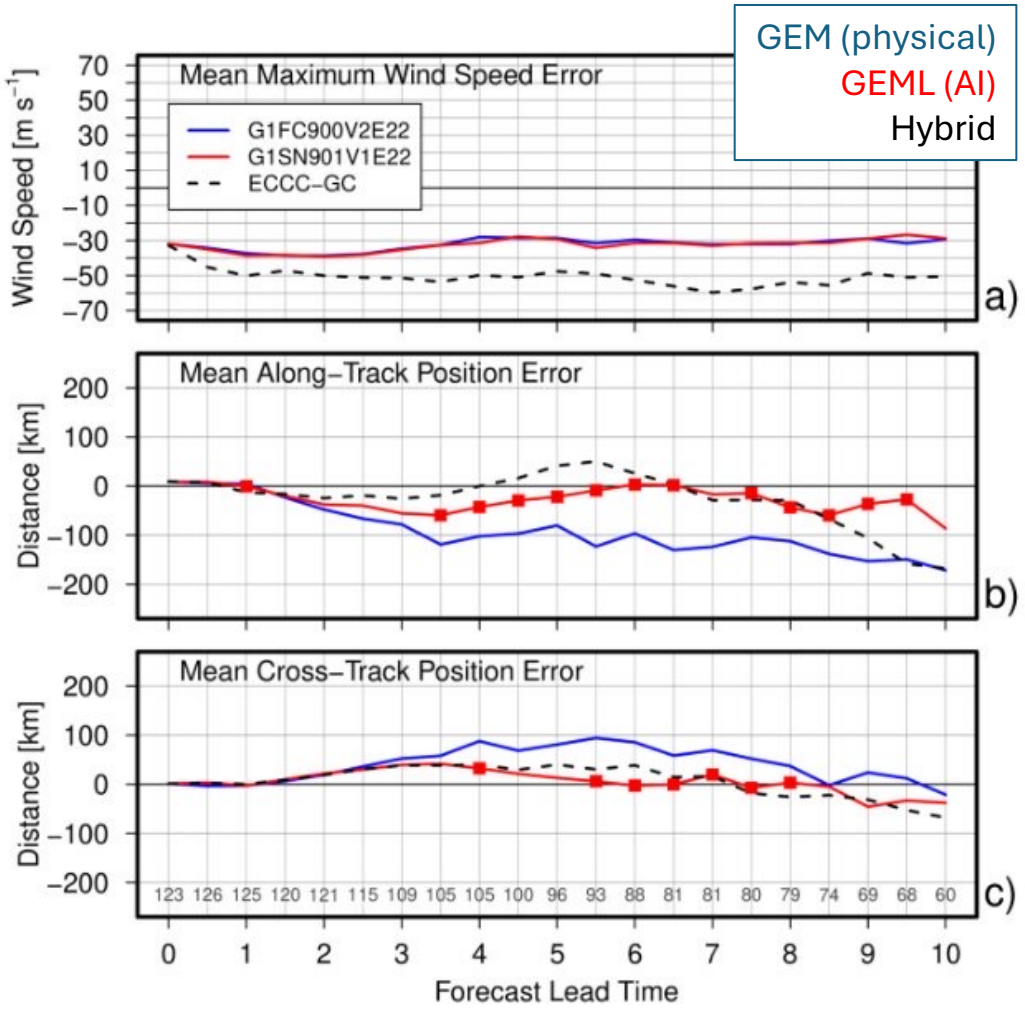

Global tropical cyclone biases wrt IBTrACS.

GEML suffers from a large weak-intensity bias, which is corrected to match the physical model in the hybrid system.

The TC environment steering flow is improved by AI nudging:

- GEM slow bias is corrected
- GEM right-of-track bias is corrected

The hybrid model is the best of both worlds for tropical cyclones: physical-model intensity with AI-model tracks.

Retraining GEML

- The GEML rollout provides the large-scale state for spectral nudging:
  - Google DeepMind GraphCast architecture
  - Improvements in the AI model directly benefit the operational hybrid system
- Licensing requires from-scratch retraining.
- Trained on ERA5 (1979-2016) and fine-tuned with IFS HRES analyses through 2021.
- Retrained GEML equivalent or superior to original GraphCast via bias/std dev tradeoffs (optimize to a different local minimum).

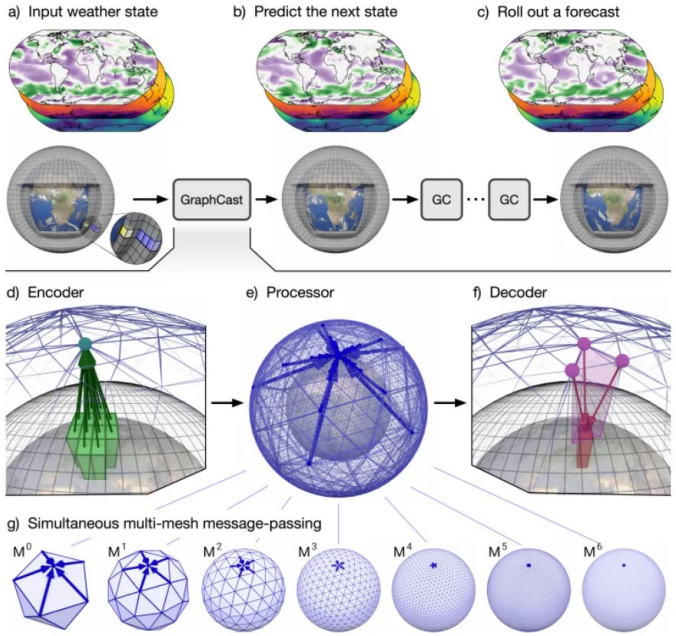

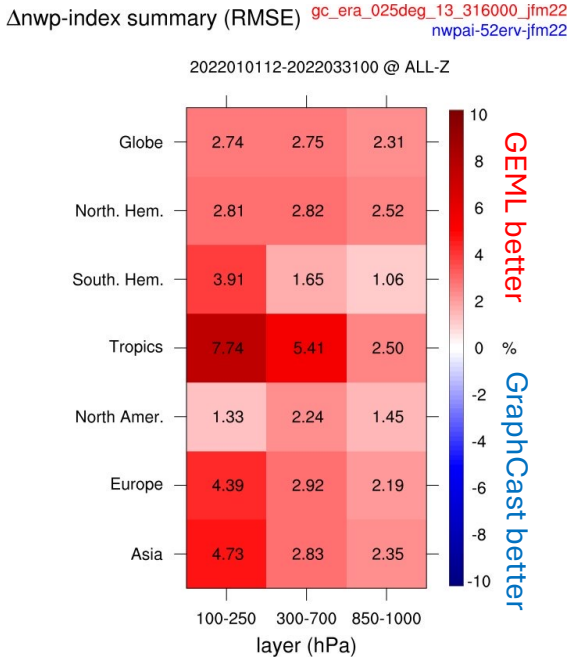

GraphCast schematic (top; Lam et al. (2022); arXiv). Difference in forecast skill index between GraphCast and GEML for the winter 2022/23 period (left).

PARADIS Development

ECCC is developing the PARADIS (Physically-Insipred Advection, Reaction and Diffusion on the Sphere) AI model

Architecture is motivation by PDE principles through hard physical constraints

PARADIS has a convolutional neural network backbone

Stacked layers emulate PDE splitting, combining a semi-Lagrangian advection step with reaction-diffusion modules

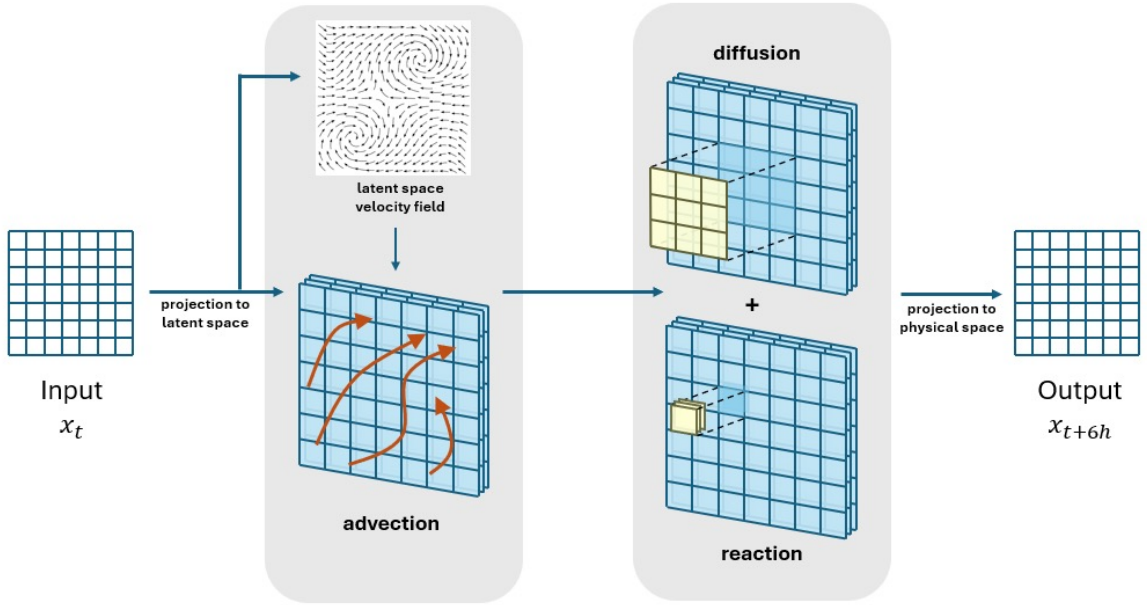

Schematic of PARADIS architecture (courtesy of Carlos Pereira Frontado).

PARADIS

PARADIS Progress in 2025

- **Architecture design:**
  Physically-inspired advection–diffusion–reaction framework under active development, with multiple architectural variants being evaluated.
- **Low-resolution experiments:**
  Extensive training and ablation studies completed at coarse resolution to validate design choices, ensure numerical stability, and optimize training efficiency.
- **Medium-resolution (1°) results:**
  Training demonstrates strong short-lead forecast skill and physically consistent dynamics.
  Work is ongoing to reduce error growth and grid-scale noise at longer lead times.
- **Upcoming work:**
  Transition to **¼° training** on multi-GPU clusters planned for late 2025, focusing on scalability, higher-resolution dynamics, and operational feasibility.
- **Next steps:**
  - Benchmark against baseline models (e.g., GraphCast-like) using external and internal analyses.
  - Conduct systematic long-lead validation and diagnostic analyses.

Computing Upgrade in 2026

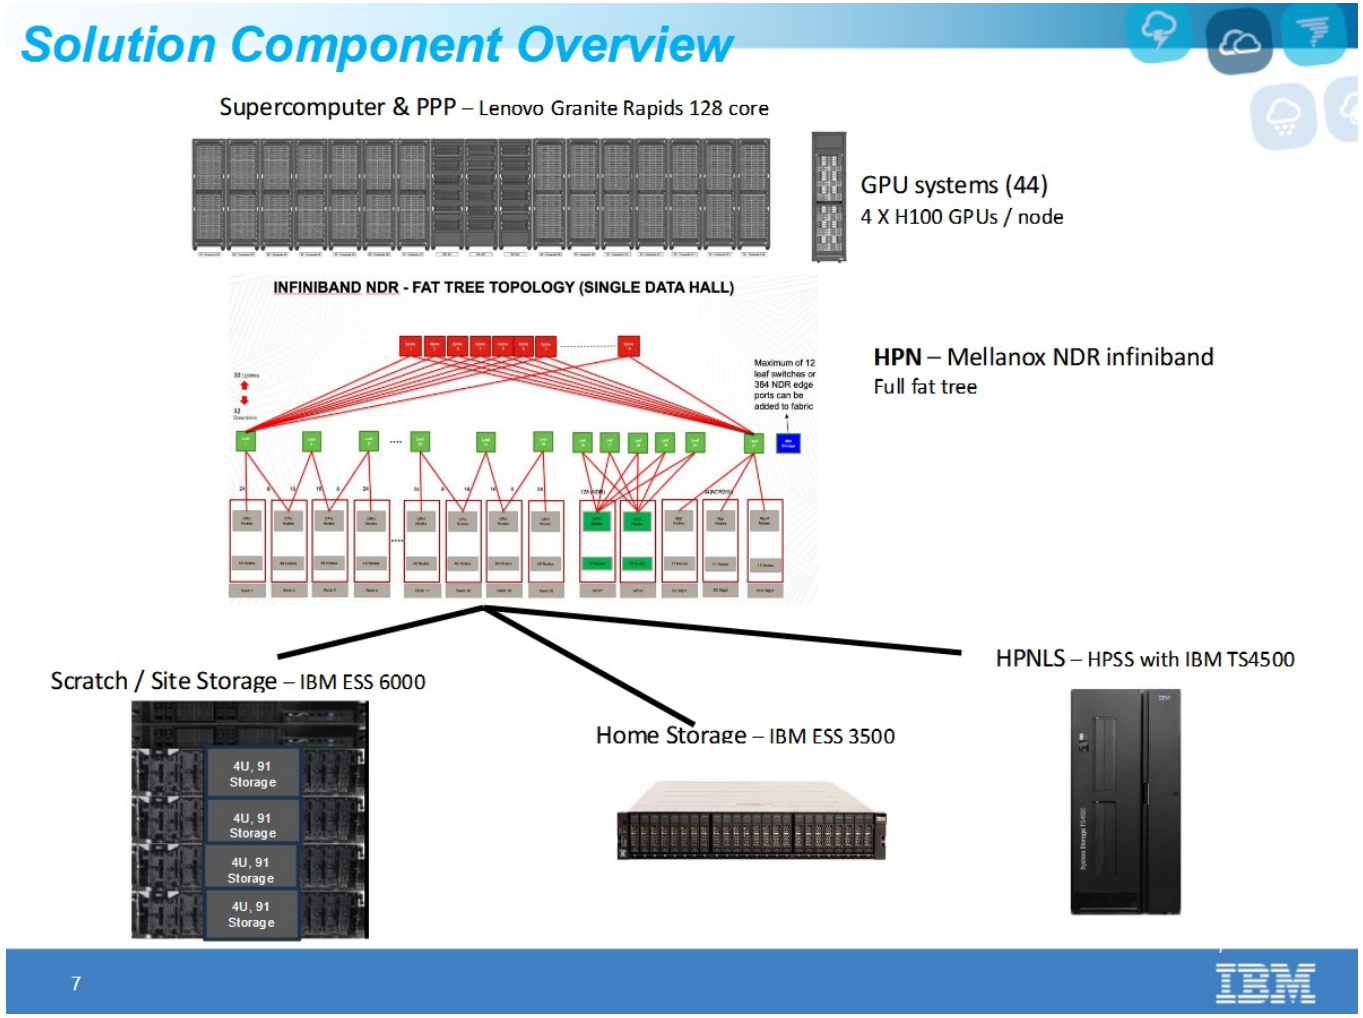

ECCC is currently completing an upgrade to our computing infrastructure:

- 239K to 377K cores
- Cores/node increase from 80 to 256
- Memory drops from 6 GB to 3 GB per core (large-memory nodes for DA)

The GPU count increases to facilitate AI model development and training, and GPU porting of existing codes:

- 2xA100 nodes (4 GPU/node) being upgraded to 44xH100 nodes (4 GPU/node)

Ongoing Physical Model Development

The ECCC physical model (GEM) is still used in all NWP applications, and dycore supports climate projections.

In an operational-AI future, physical models are needed to generate training unless direct-from-observations successful.

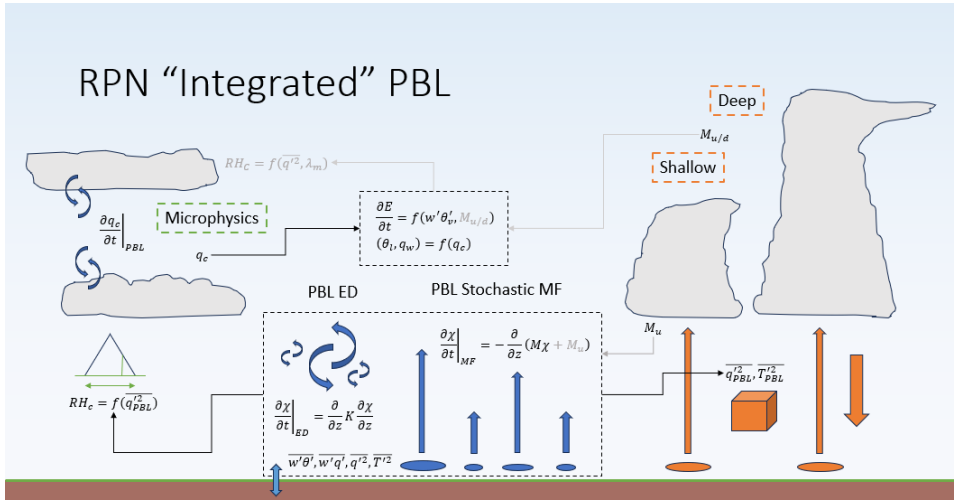

The new “RPN Integrated” PBL scheme is a 1.5 order closure that is tightly coupled to the model microphysics scheme to create and PBL clouds:

- RPNInt eliminates longstanding biases in marine Sc regions.

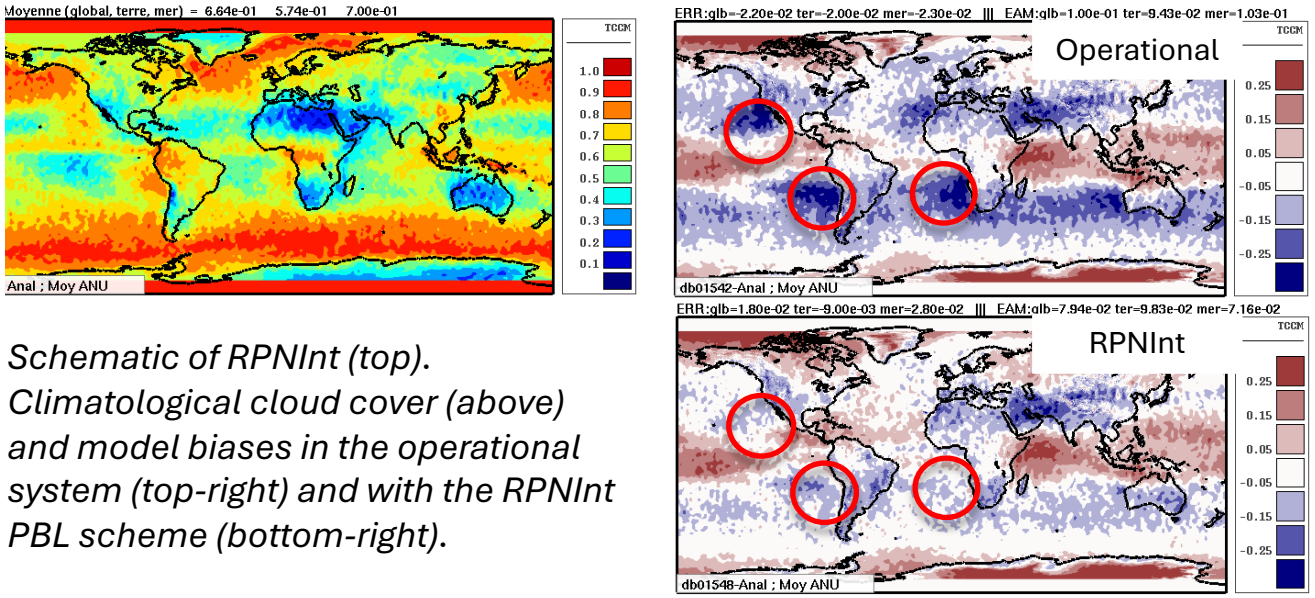

Schematic of RPNInt (top). Climatological cloud cover (above) and model biases in the operational system (top-right) and with the RPNInt PBL scheme (bottom-right).

Operational Upgrade in 2027

- Innovation cycles at ECCC are set at 30 months by supercomputer upgrades – very long compared to other centres
  - Increases turnaround time for forecaster feedback-based corrections
  - Makes each cycle very large, including 40 individual systems
- Work towards the 2027 upgrade needs to be finished in early 2026 for integration testing:
  - Ensemble DA (256) and forecast (20+1) to move from 25 km to 15 km grid
  - All systems to move from pressure-based to height-based dynamical core for improved stability in steep terrain
  - Major upgrade planned to model physics for the 2.5 km High Resolution Deterministic Prediction System
- Details and results for WGNE 2026!