

Numerical
Model Training Course



2020

ICON Training in the Cloud Lessons Learned

F. Prill and C. Eser (DWD) | Weather and climate in the cloud | 2021-02-08

ICON Training Course 2020

- Between November 16 and 27, carried out virtually for the first time
- 40 participants from home and abroad

Contents of the course:

Numerics, physical parameterizations, technical aspects of the ICON numerical weather prediction code (in total 16x1h lectures)

4 x 2h practical exercises with “teams” of two participants each.

- Use of the community cloud **"European Weather Cloud"** as **laaS** (computing and storage):

Classical example for dynamic requirements with short-lived containers, realized by “the cloud”.

ICON in the Cloud?



ICON = ICOsahedral Nonhydrostatic Model

Joint development project of **DWD** with **MPI-M**,
the **German Climate Computing Centre** and **KIT**.

*finite volume NWP model with icosahedral-triangular grid,
grid nesting, 2nd order DyCore.*

Made available to the scientific community since 2015.

NHMSs Capacity Building

2021: World Bank-financed central Asian PoC

... investment in improved weather services.

- Do cloud services and containerization offer a new way of deploying ICON to regional MetServices and universities?



Pilot domain: mountainous area,
3.29 km horizontal mesh size

Containerization Concept

All major cloud vendors offer container support (but can also be used on-premise): Modular principle

- Containers are **not virtual machines**, but run as processes like normal software on the system.
- Can be used and reproduced almost anywhere. No need to consider **software dependencies** or library versions.
- BTW: Provides a formal **recipe for the installation and usage** of the ICON model and its related tools.

ICON-LAM

Software Libraries



Host OS

Server (Machine)




icon-pre

pre-processing of input data from DWD's operational database and other sources

icon-post

post-processing and basic visualization of ICON runs

– *single instance, I/O intensive* –



icon-lam

self-contained, MPI-parallel executable of the ICON model

icon-debug

source code and compilation environment

– *CPU + network intensive* –


Capacity Building and Training



DWD has organized annual NWP training since the early 2000's.



- targeted at national MetServices and scientific users who are planning to run DWD's regional models.
- 30 – 60 participants each year.
- cooperation with MeteoSwiss, the CLM Community, NMA Romania, Karlsruhe Institute of Technology.



The 2020 Virtual Course

- Lectures were recorded; available as browser stream after the course.
- Course materials: ICON manual with exercises, ICON source code and test data.
- Practical exercises with “teams” of two participants each.




EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED
ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

Numerical Model Training 2020: Program Overview						
Lectures	Monday 16.11.	Tuesday 17.11.	Wednesday 18.11.	Thursday 19.11.	Friday 20.11.	
	10:30 – 11:30	Introduction <i>D. Rieger</i>	Dynamics <i>G. Zängl</i>		Physics Overview <i>D. Klocke</i>	Turbulent Diffusion <i>M. Raschendorfer</i>
Exercises	11:30 – 12:30	Getting Started I <i>F. Prill</i>	Getting Started II <i>F. Prill</i>		Advection <i>D. Reinert</i>	Turbulent Transfer <i>M. Raschendorfer</i>
	12:30 – 14:00	Lunch Break	Lunch Break		Lunch Break	Lunch Break
	14:00 – 15:00	Exercise I	Exercise I		Exercise II	Exercise II
	15:00 – 16:00	Exercise I	Exercise I		Exercise II	Exercise II
Week 2						
	Monday 23.11.	Tuesday 24.11.	Wednesday 25.11.	Thursday 26.11.	Friday 27.11.	
	10:30 – 11:30	Microphysics <i>A. Seifert</i>	Gravity Wave Drag <i>M. Köhler</i>		Lake & Sea Ice <i>D. Mironov</i>	CLM <i>C. Steger</i>



Each exercise consisted of progressively complex building blocks.

- **ICON-LAM (limited area model) run** driven by DWD operational data.
- **Pre-processing** of initial and boundary data for a computational domain over Germany.
- **Visualization** of regional model output.
- Studies: temporal resolution of boundary data.
- Modification of the **model code**, implementation of an own diagnostic.



DWD training course team:
D. Reinert, D. Rieger,
FP, CE

Containers are **immutable**: all state information is loaded during startup.

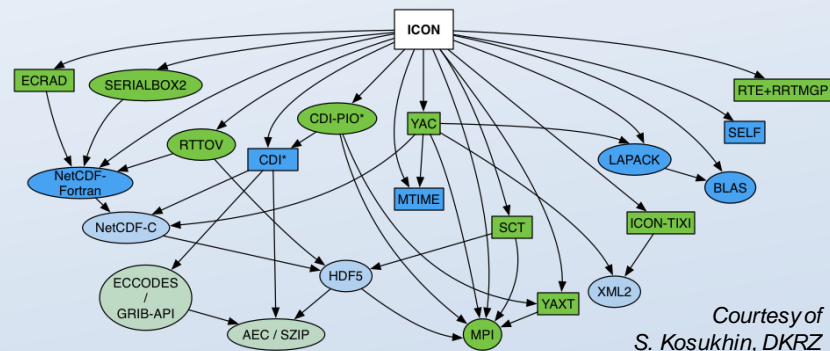
No expertise necessary to set up **icon-lab**, access via web browser.

ICON-LAM executable +

- Jupyter Notebook ("web terminal")
- OpenMPI runtime environment
- I/O libraries: NetCDF, HDF5, ECCodes (preset for DWD)
- Python + NCAR's PyNGL library
- CDO tools (climate data operators)
- `gcc` development tools

copy-in of large data sets:

EWC S3 storage bucket used for grid data, etc.

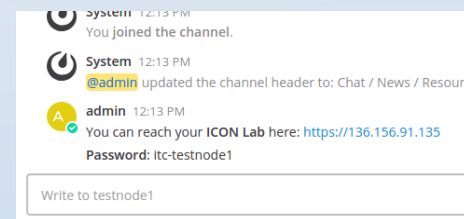
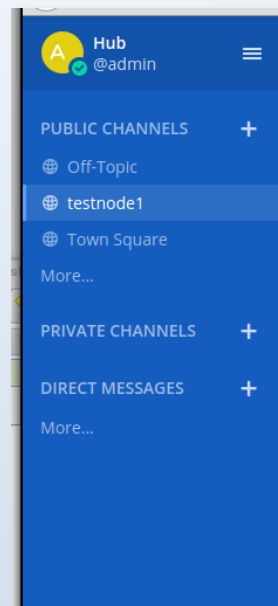


Close View: Training Hub

The whole course was organized and held together by the **“training hub”**.

Mattermost team communication software running for the whole course duration.

- Chat platform container with sub-channels and integrated video chats.
- Automatized messages: Automatic creation of channels and “bot messages”
- Custom setting of the JupyterLab credentials posted automatically in the chat.



EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE, FOCUSED
ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

Controlling the EWC Instances

The icon-lab containers are set up by **Morpheus App Blueprints**.

Additionally: **interactive “icon-operating” container** executed locally (at DWD) offering convenience functions.

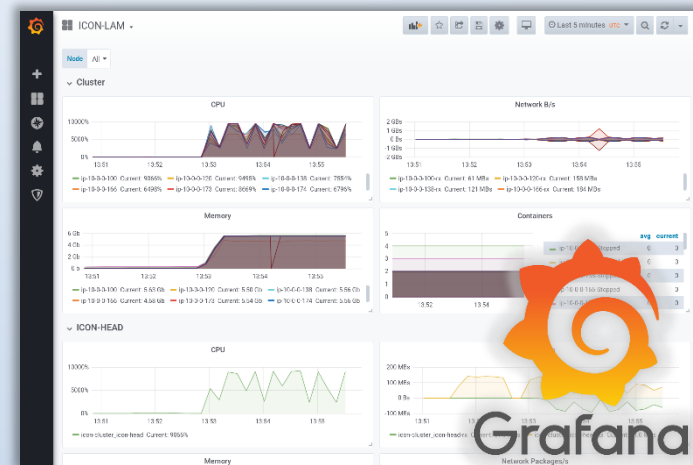
- retrieve IP addresses, execute scripts
- automatically post chat messages



Behind the scenes:

- ICON-LAM containers include remote monitoring via web browser (Grafana/InfluxDB/Telegraf).

CPU load, memory consumption, network bandwidth, network packets, ...



November 16 - 27, 2020

Lessons Learned



Feedback of the participants mostly very positive.

As always,
there is room for improvement...

- Time slots for the exercises (2 weeks) could be longer.
- Some exercises could be simpler; walkthroughs.
- The technical setup could be supplemented by GoogleDocs-like, real-time web documents.
- November 16: A problem with S3 had us panicking for a bit (however, resolved by the EWC team!)



The **icon-hub** successfully unified communication and “bot messages”.

- You really need a background channel via which the teachers can coordinate. Attachments (timetables etc) sometimes drown in the stream of other posts.

The **icon-lab** proved to be a ready-to-run sandbox.

- 8 CPUs essential (ICON wouldn't work with less).
- It is a good idea to separate the course material (as a copy-in tar.gz file) from the container image.

Features that were missing this time:

- **Real time collaboration** still primitive in JupyterLab. Next time, we'll use JupyterLab RTC (unstable).
- Registry for Container Images and DNS (Certificates).



Virtualizing the ICON course cost quite a bit of preparation.

- We entered new territory here, but everything went well.
- Impossible without a whole team of DWD colleagues. Essential support through the whole ICON working group and the EWC team!
- With the introduction of Kubernetes in the EWC, we will switch to it in the future.

Cloud technology: **scales dynamically**

... but, for a training course, this is not true:
the teachers did not scale :)

ICON-LAM

Software Libraries

K8s Container
Runtime Interface



Host OS

Server (Machine)