Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Big data in critical infrastructure: Production and failover infrastructure in DWD's central data management

Daniel Lee, German Weather Service (DWD)

Agenda

1. DWD's goals
2. Operational systems
3. Technical infrastructure
4. Current challenges
5. Future plans

What does DWD do?

The DWD law

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

- Monitoring of meteorological interactions between the atmosphere and other environmental systems
- Prediction of meteorological events
- Monitoring and prediction of the movements of radioactive trace particles
- Operation of the necessary observation systems
- Storage, archival and documentation of meteorological data and products

Target audiences

- DWD aids in protecting lives and property, as well as in planning and maintaining critical infrastructure in the areas of:
  - Aviation
  - Seafaring
  - Agriculture
  - Energy
  - Climate
  - Weather warnings
  - Protection and recovery from high impact weather
  - Etc.

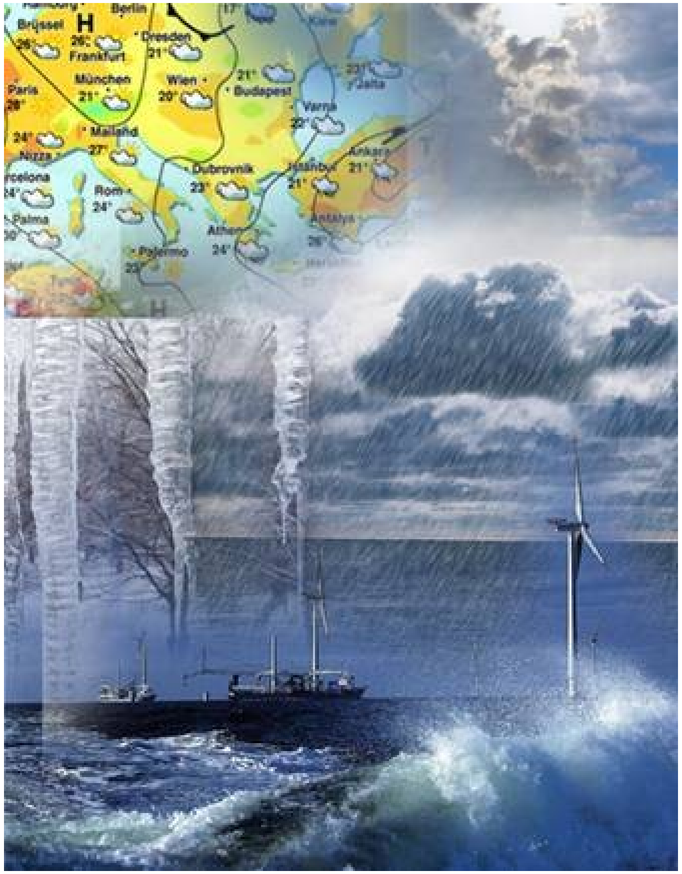

Multiscale topics

Future

Past

100 years / 10 years	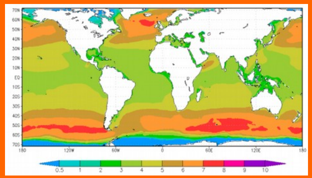	• Climate projection
1 year / 1 month		• Yearly, seasonal and monthly forecasts
today	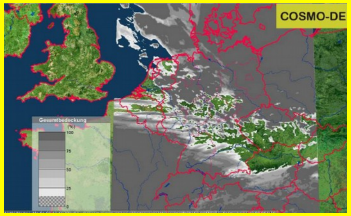	• Longterm forecasts (72-360 hours) • Midterm forecasts (12-72 hours) • Shorterm forecasts (2-12 hours) • Nowcasting (< 2 hours)
	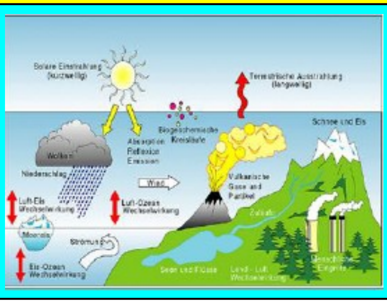	• Weather and climate observation

Example: Model output

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

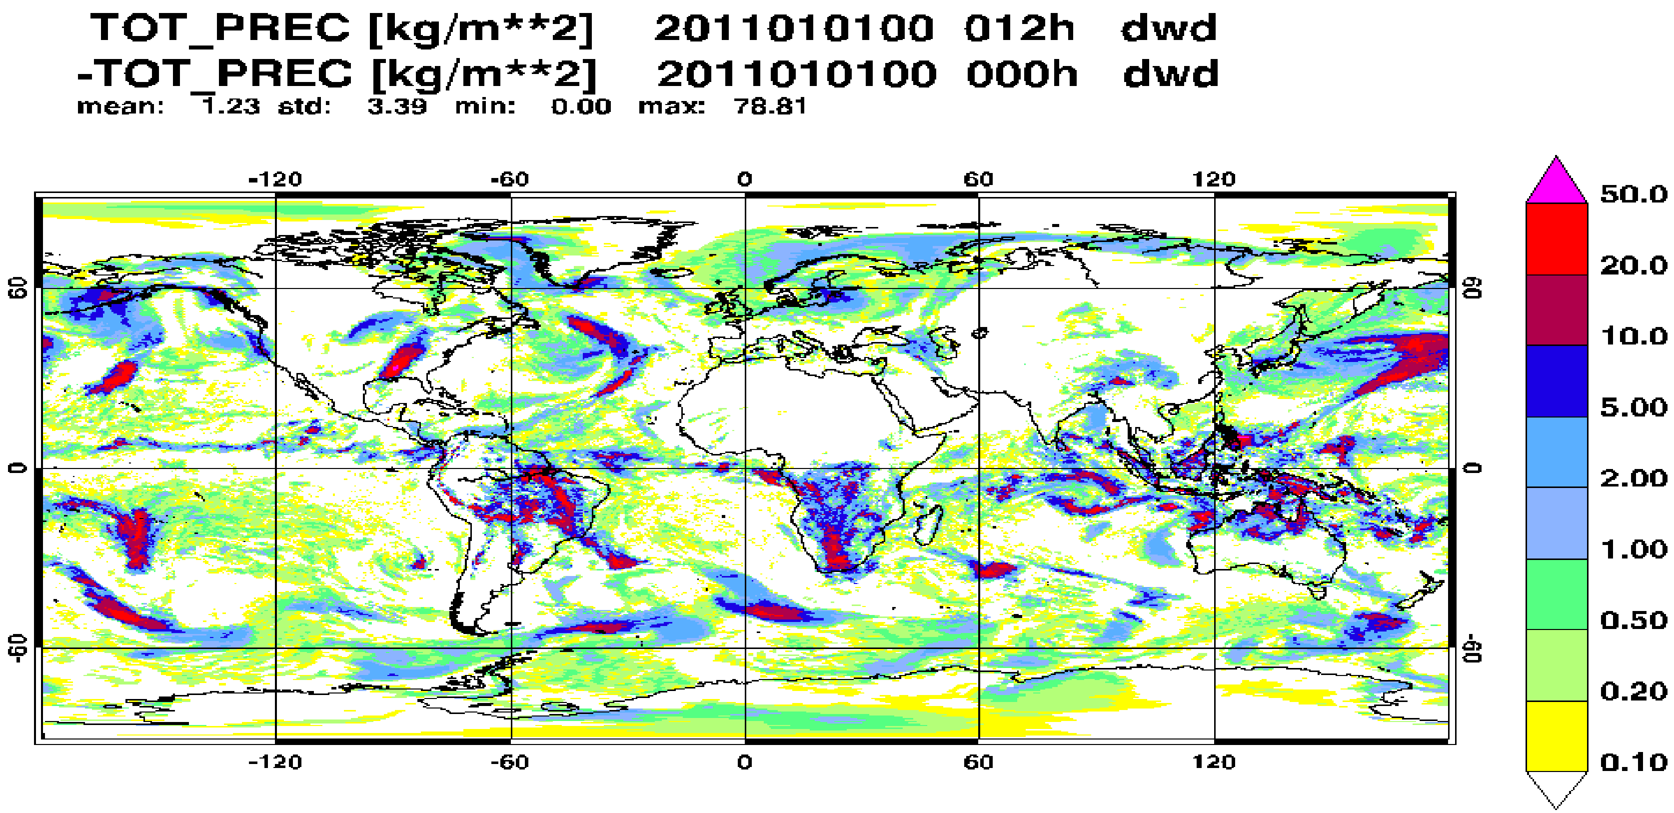

Example: Radar map

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

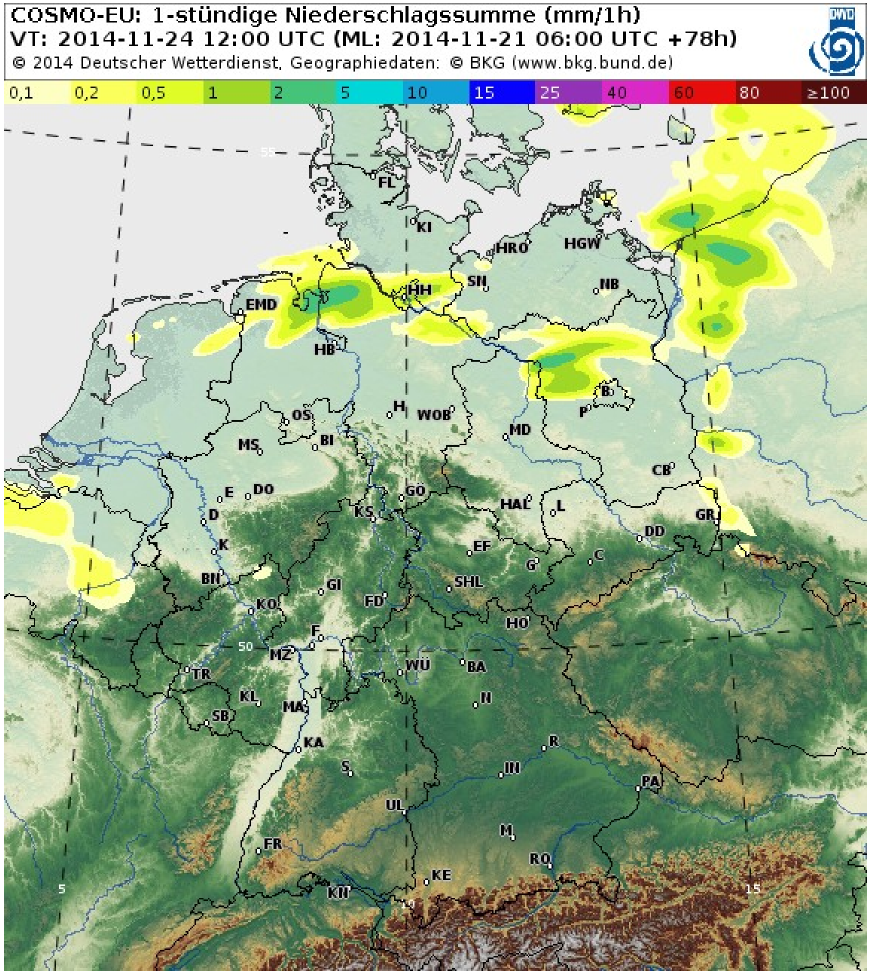

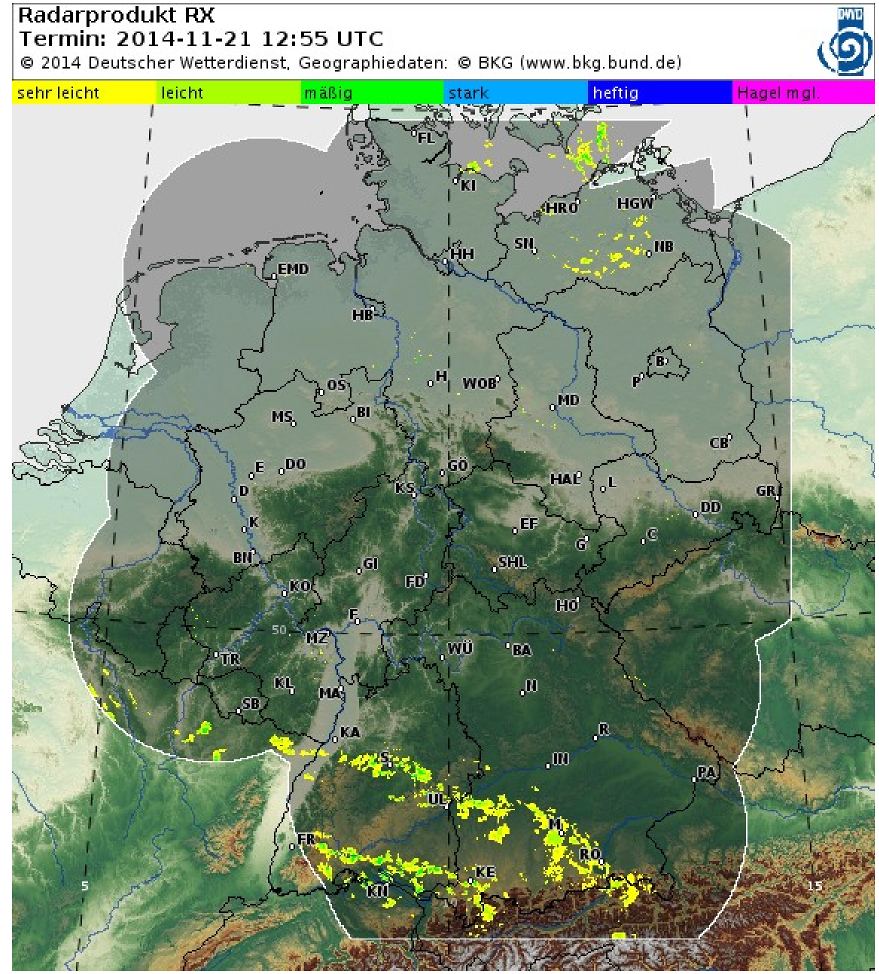

Example:
Short-term storm cell forecast

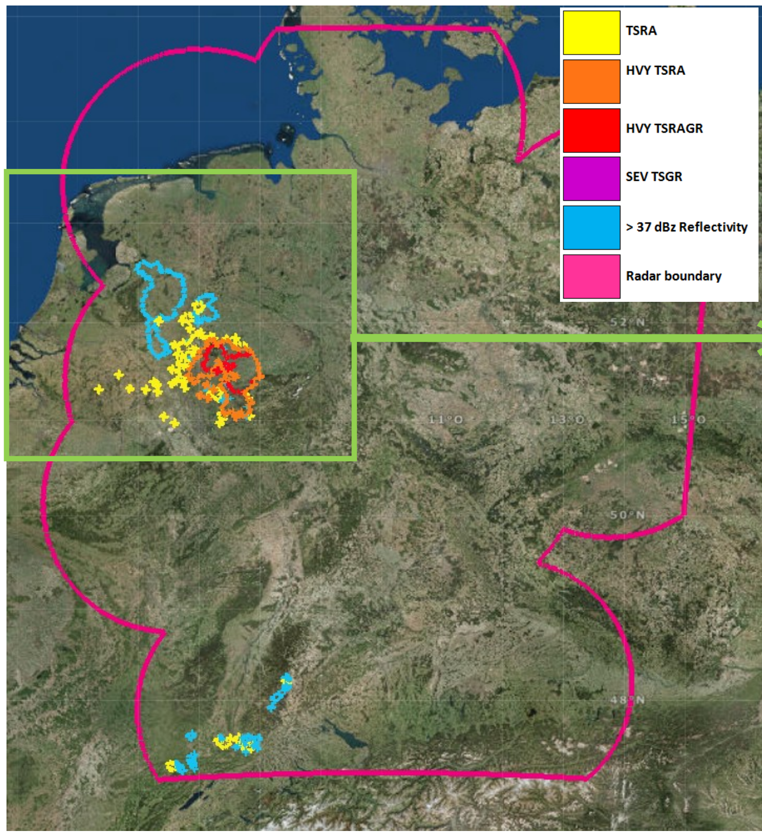

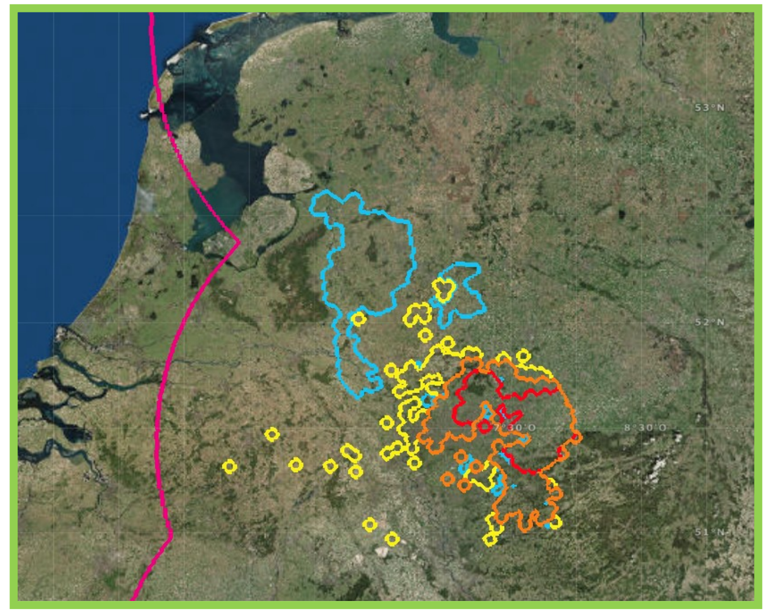

Example:
Ensemble probabilities forecast

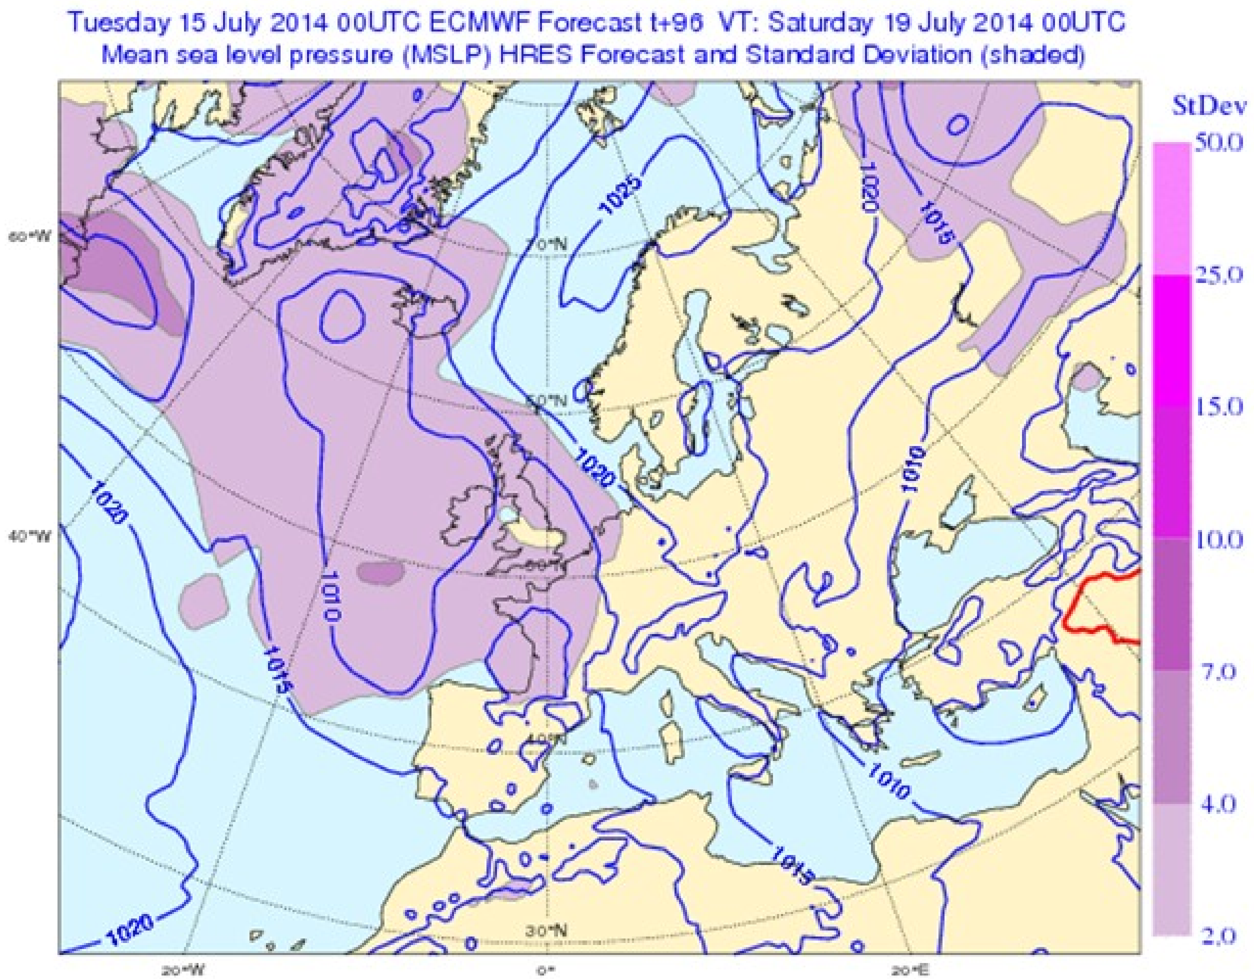

Example:
Flight cross-section

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

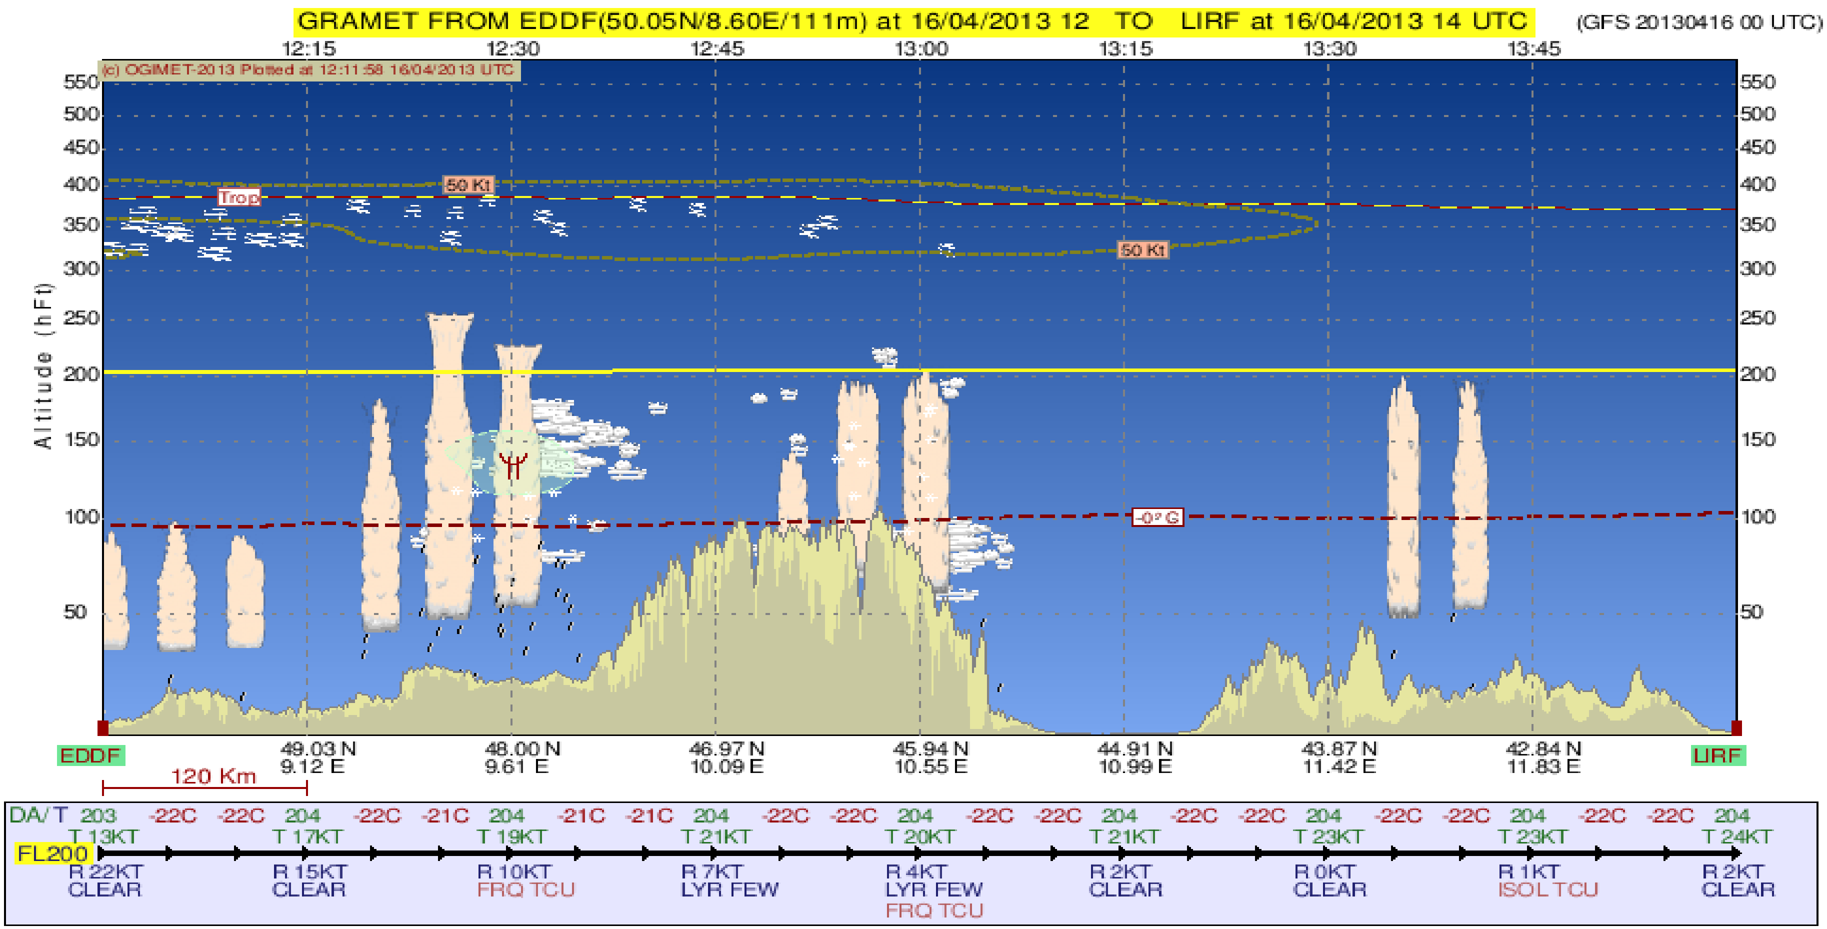

Example storm event

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Example storm event

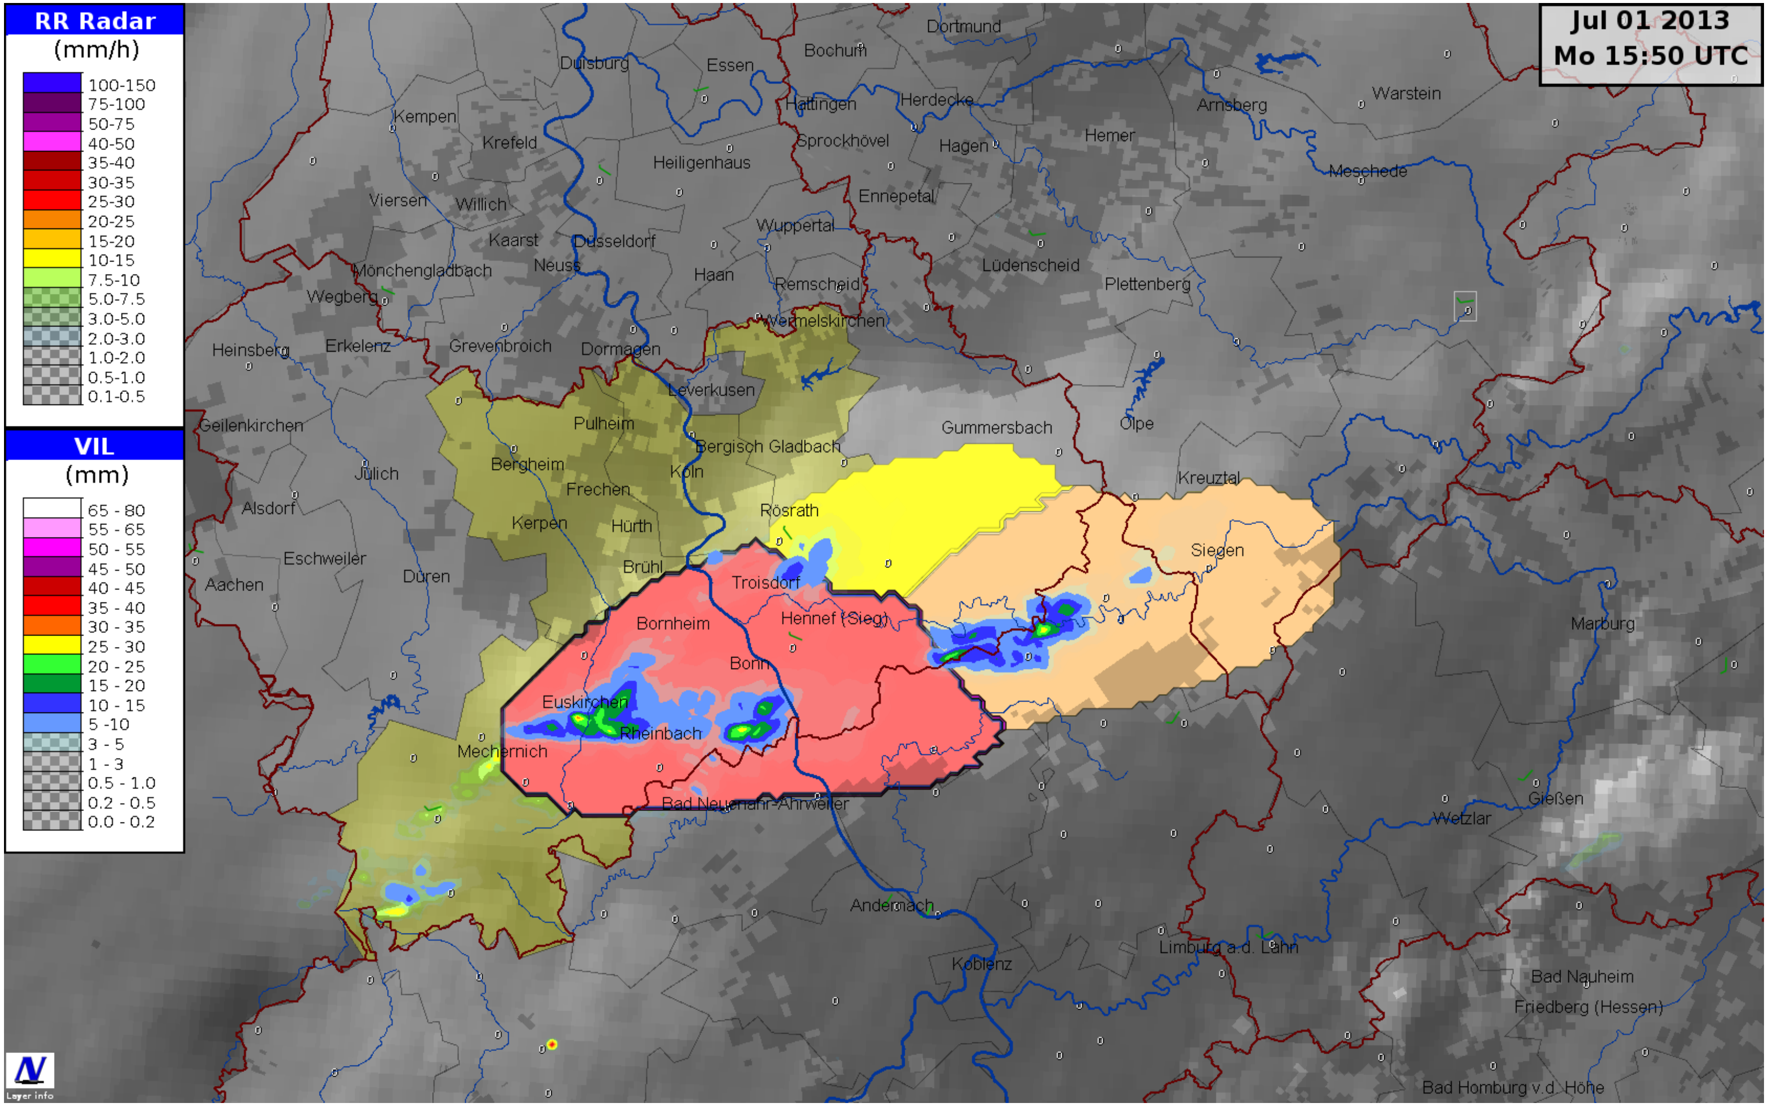

Automated product generation

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

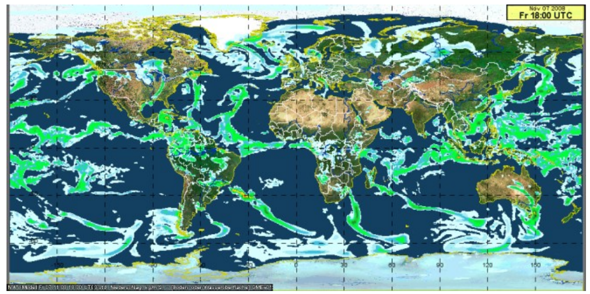

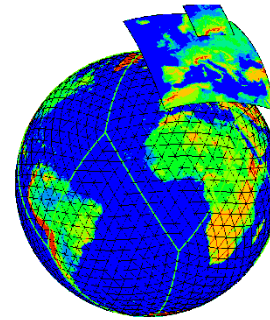

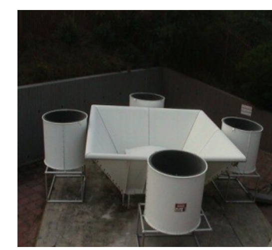

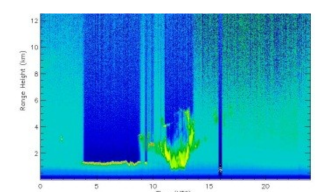

Operational systems

Models for multiple spatiotemporal scales

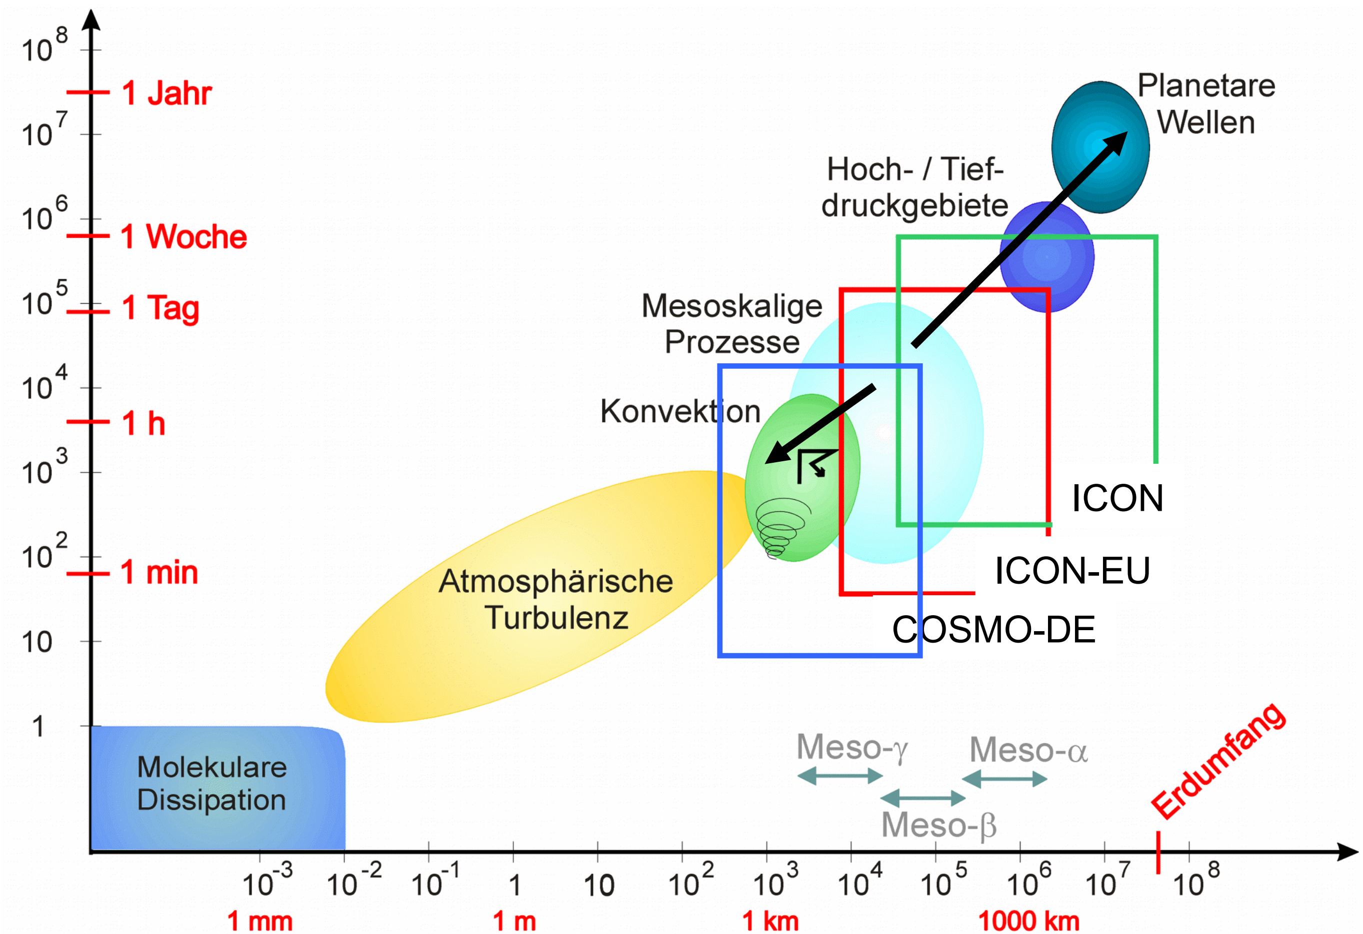

Physical observation system

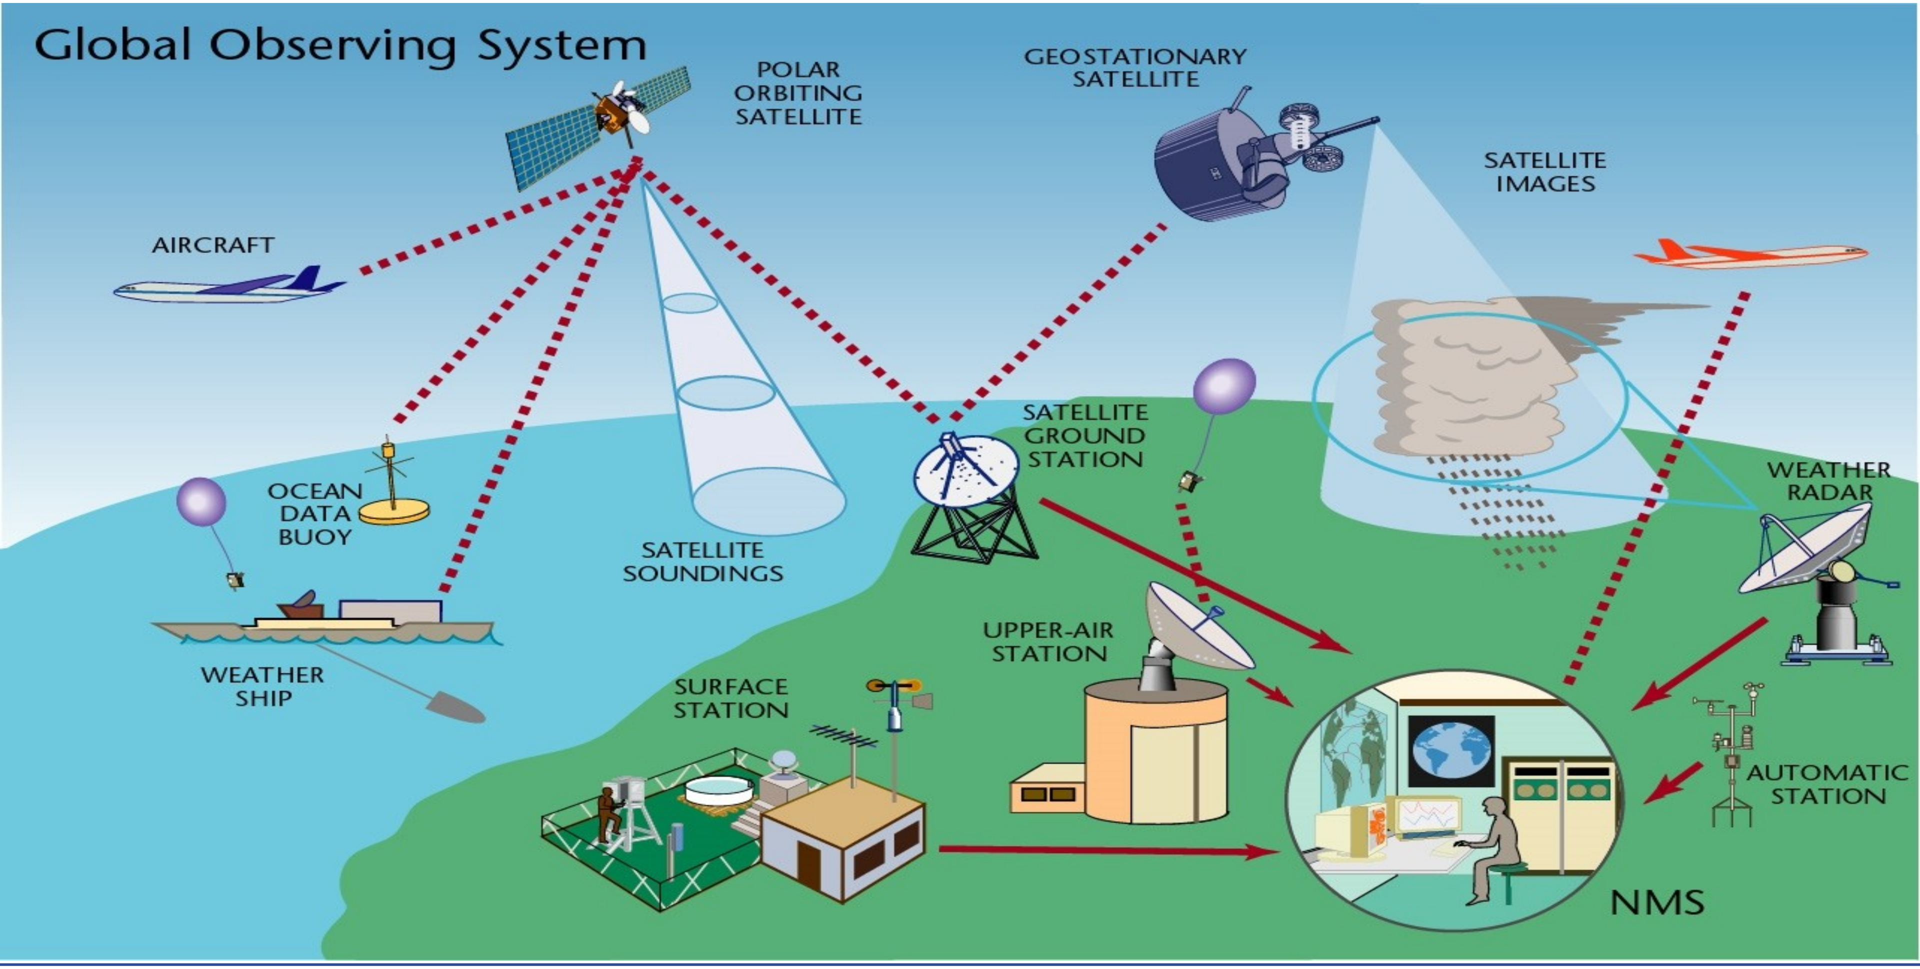

Bringing the data together

DWD
Deutscher Wetterdienst
Wetter und Klima aus einer Hand

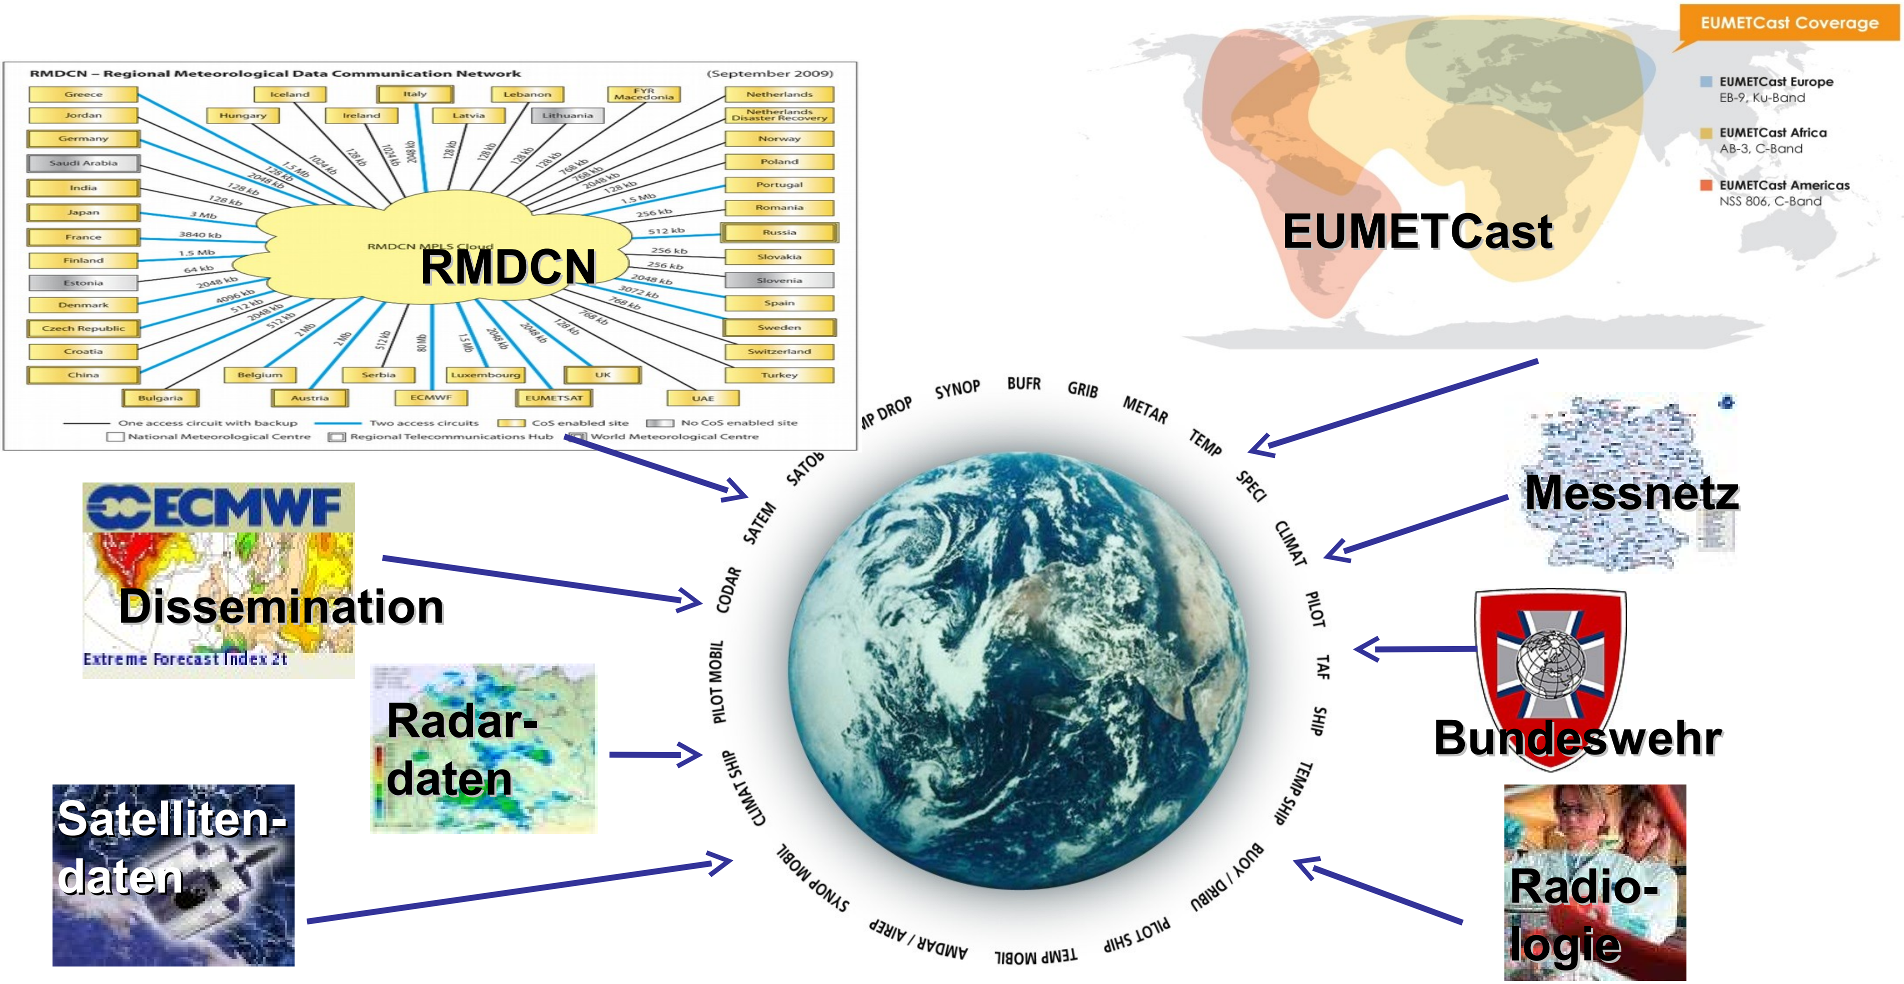

DWD's weather models

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

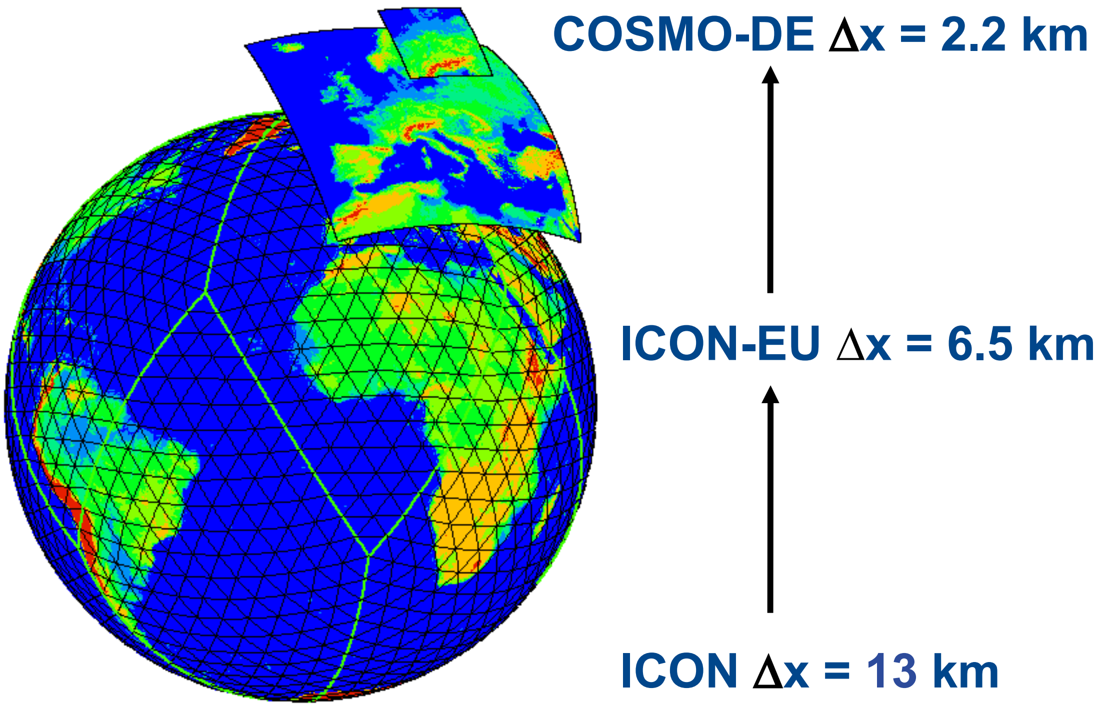

COSMO-DE (-EPS):
Grid spacing: 2.2 km
Vertical layers: 80
Forecast range: 27 hours
Runs per day: 8
EPS members: 20

ICON-EU:
Grid spacing: 6.5 km
Vertical layers: 54
Forecast range: 78 hours
Runs per day: 8

ICON:
Grid spacing: 13 km
Vertical layers: 90
Forecast range: 174 / 78 hours
Runs per day: 4

Daily deterministic output: ~ 2.5 TByte
Daily probabilistic output: ~ 3.5 TByte

Routine operation

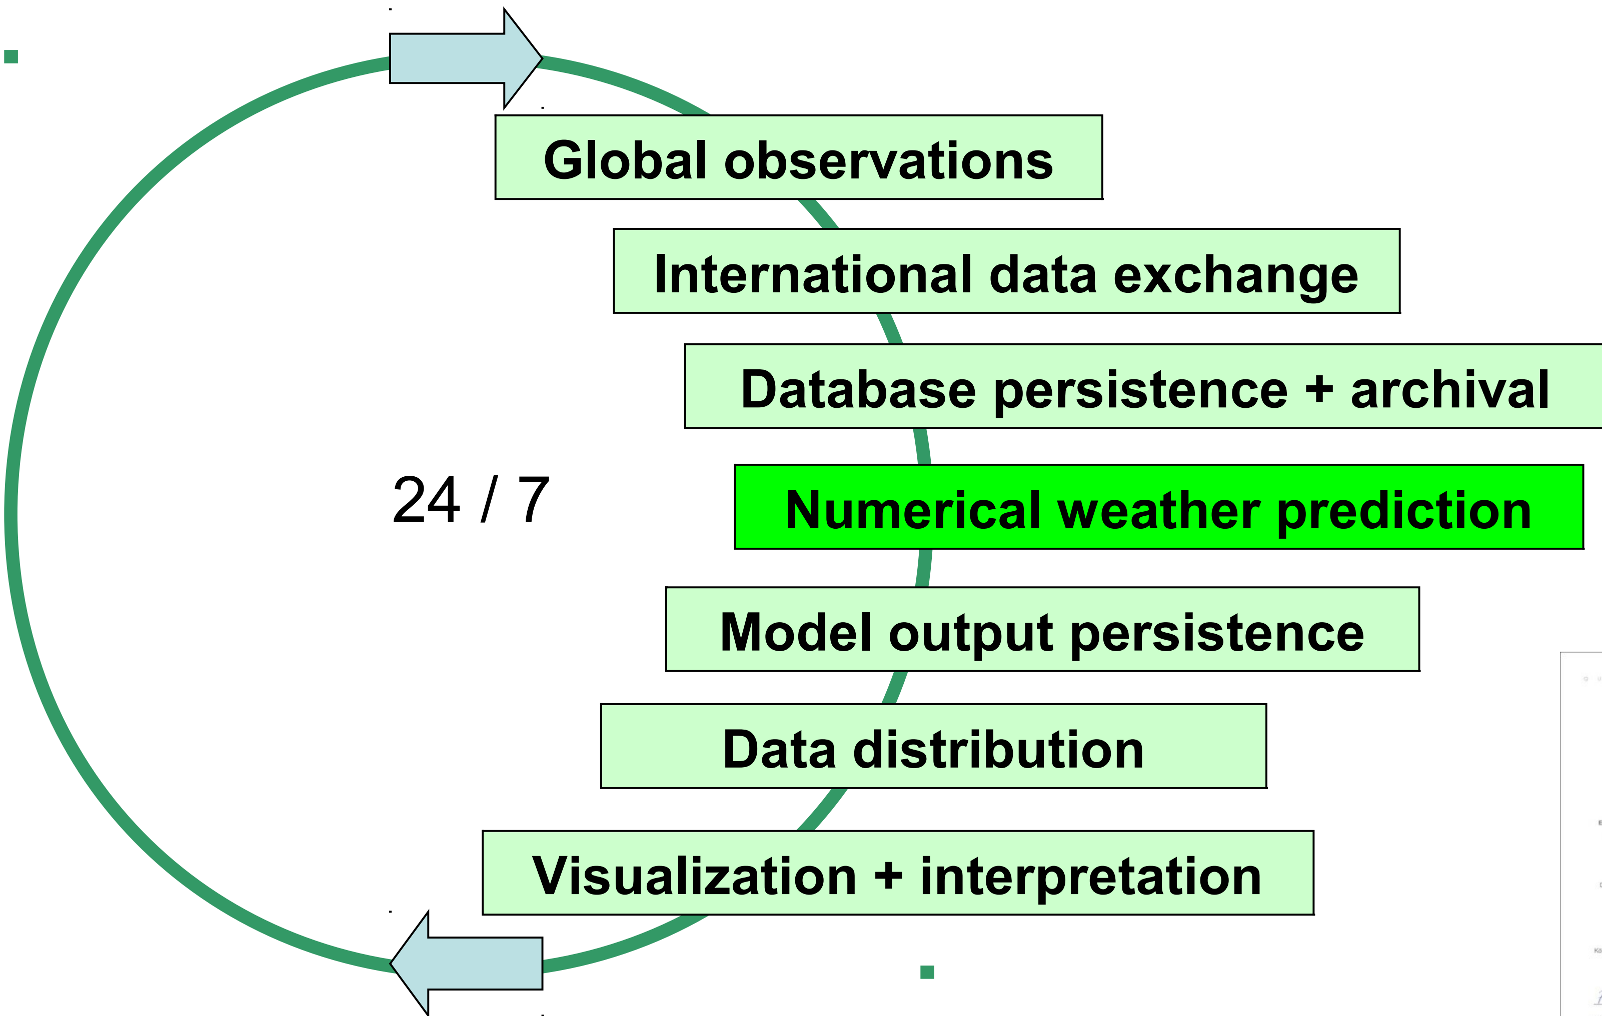


ZERTIFIKAT

Deutscher Wetterdienst
Frankfurter Str. 135
63067 Offenbach

DIN EN ISO 9001:2008

DIN EN ISO 9001:2008

Routine operation

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

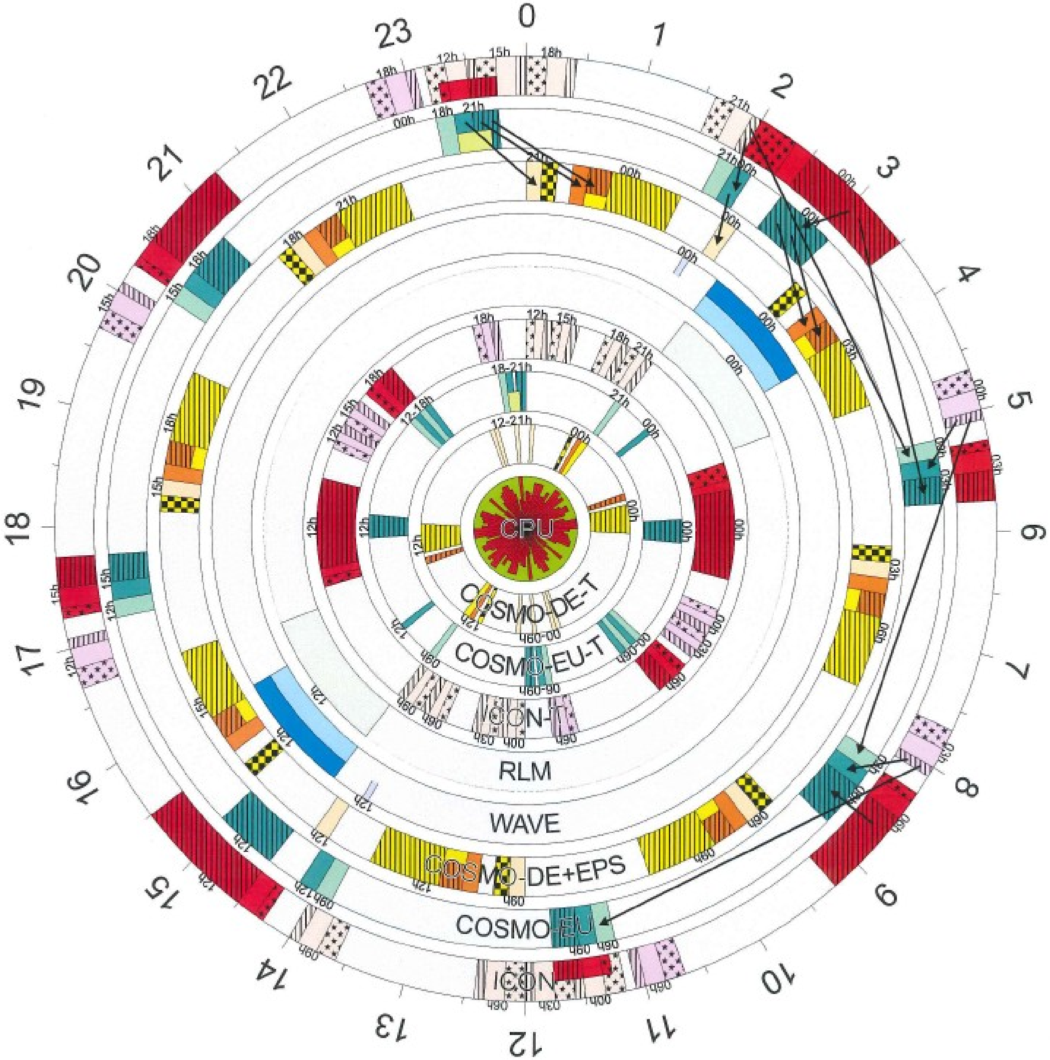

Daily data ingress / egress

Type of data	Source	Producer	# reports	Approx. daily data volume
Observation	in-situ	Manned ground station	12,000	350 GB
		Automatic ground station	50,000	
		Ship observation	2,500	
		Buoy observation	750	
		Radiosonde	900	
		Aircraft	3,000	
	remote sensing	Radar	17	
		Lidar	2	
		Wind profiler	4	
		Satellite	20	
Model output	model	COSMO-DE	8	6,000 GB
		COSMO-DE-EPS	8	
		ICON-EU	8	
		ICON	4	
		Wave model / other models	12	500 GB

Physical infrastructure

DWD's processing power

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

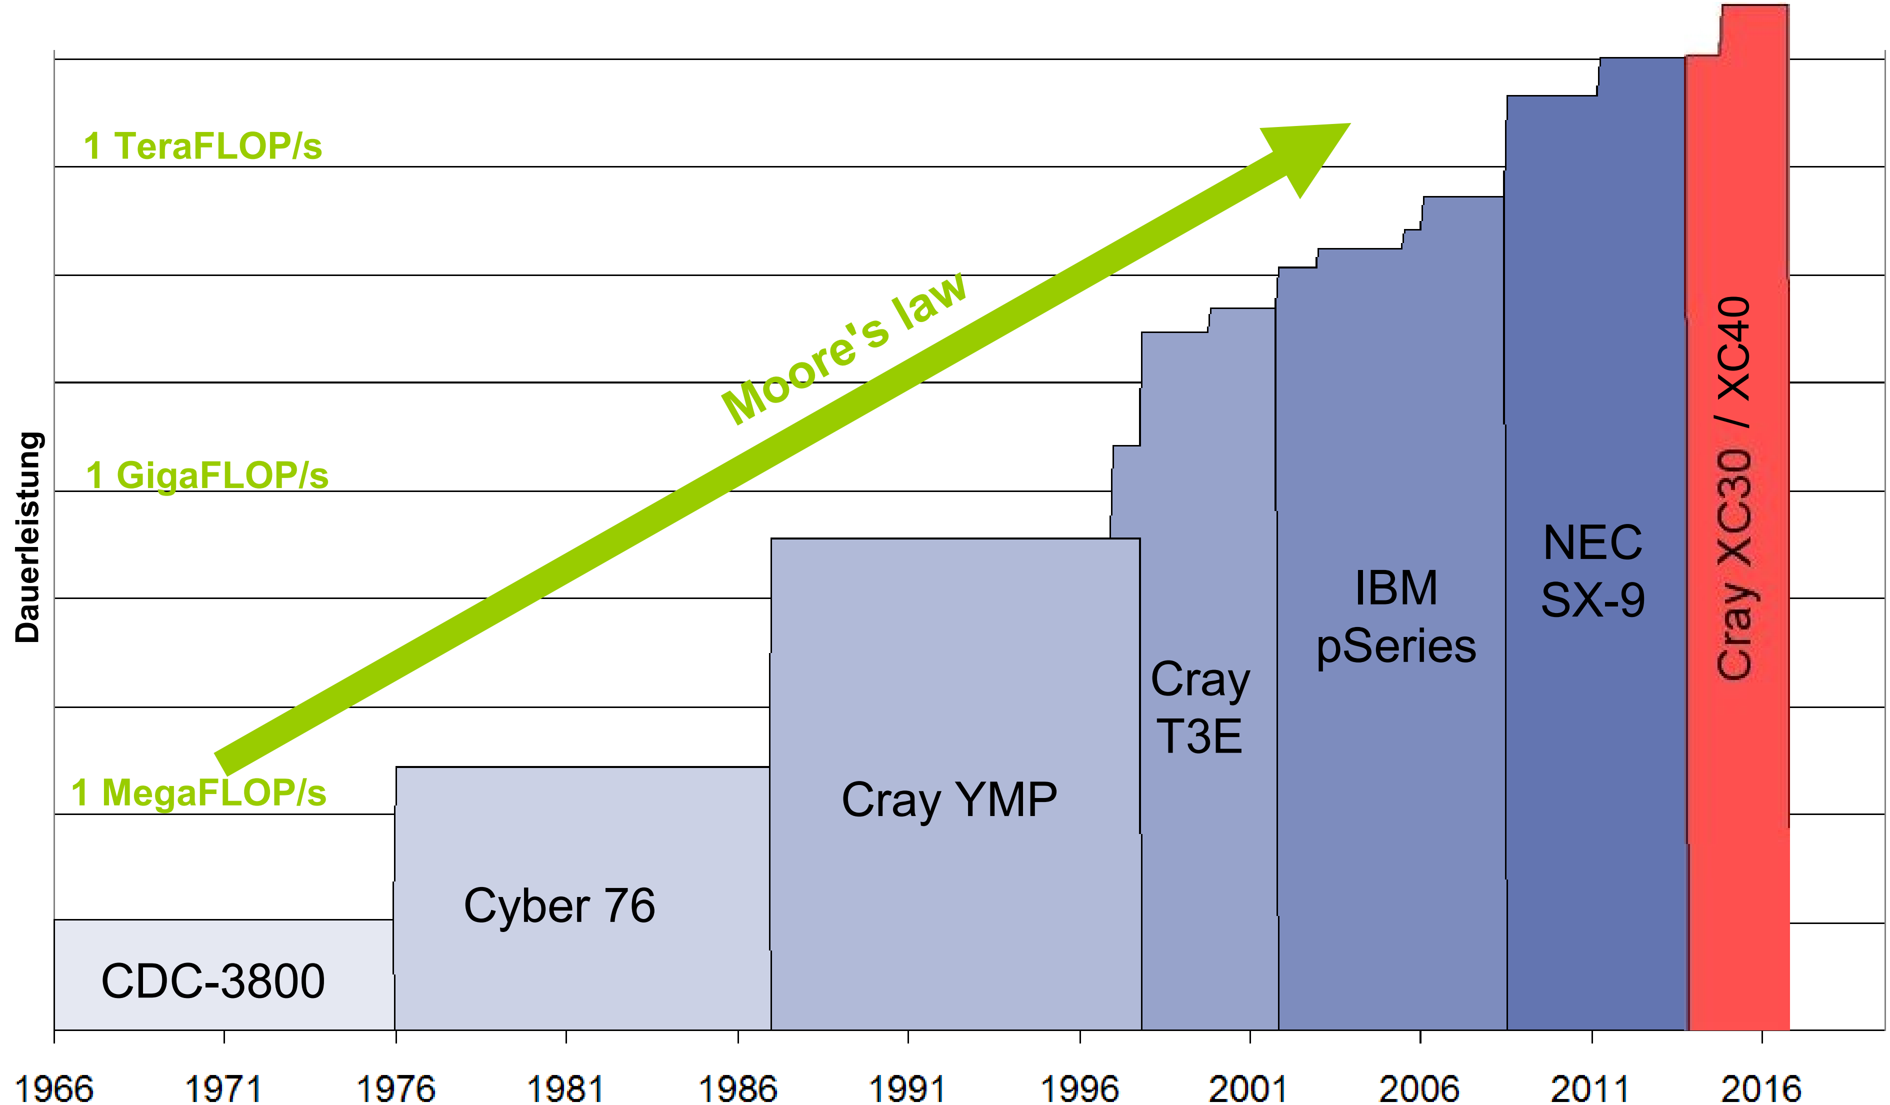

Multiple supercomputers

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Production 	Research & Development 
• 24 / 7 routine production • Data assimilation for computing model initial states • All numerical weather prediction models • Parallel routine for evaluating model changes	• R&D • NUMEX: Numerical experiments • Backup if routine computer unavailable

HPC specifications

- 2 Cray XC40 supercomputers
- Each: 796 nodes, 17,648 compute modules, 79 TB RAM
- Xeon prozessors, connected with Aries network
- Top performance: 550 TeraFLOP/s per computer

- 2 Linux Cluster (Megware Slashtwo)
- Each: 523 nodes, 4.5 TB RAM
- Each: 500 Xeon compute modules connected with Infiniband network
- Top performance: 16,7 TeraFLOP/s per computer

Volume in meteorological databases

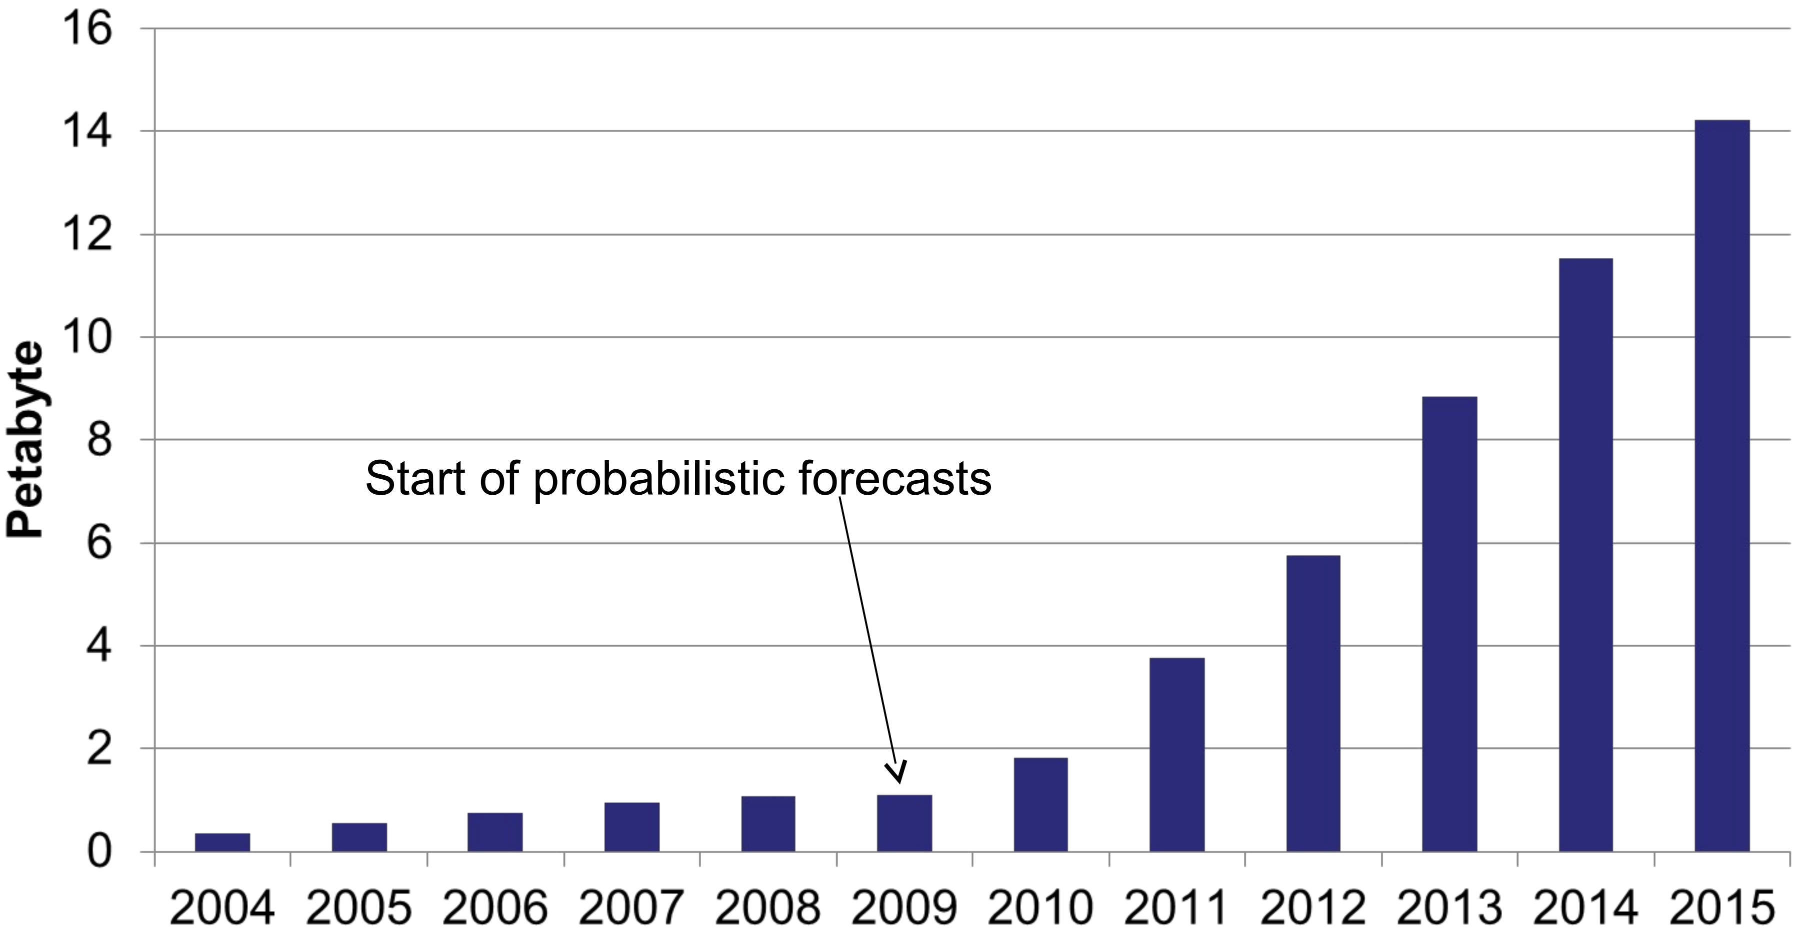

Storage systems

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

- Cray Sonexion & NEC/NetApp E5500 with ca. 4.2 PB storage capacity
- 1.2 PB global storage for HPC
- Write speed: 15 GB/s

- Also:
  - 6 data servers (NEC/SUN Fire X2-8 with 80 Intel Xeon compute modules & 1 TB RAM each)
  - 2 mirrored data servers
  - Availability: 99,9%
  - Server group manages up to 3 PB meteorological data

- Archive of 2 StorageTek SL8500 tape libraries with 10,000 tape casettes each
- > 40 casettes can be read and written to simultaneously
- 16 robots physically access the tapes and insert them into the archive server
- Estimated data volume by 2016: 60 PB

Georedundancy

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

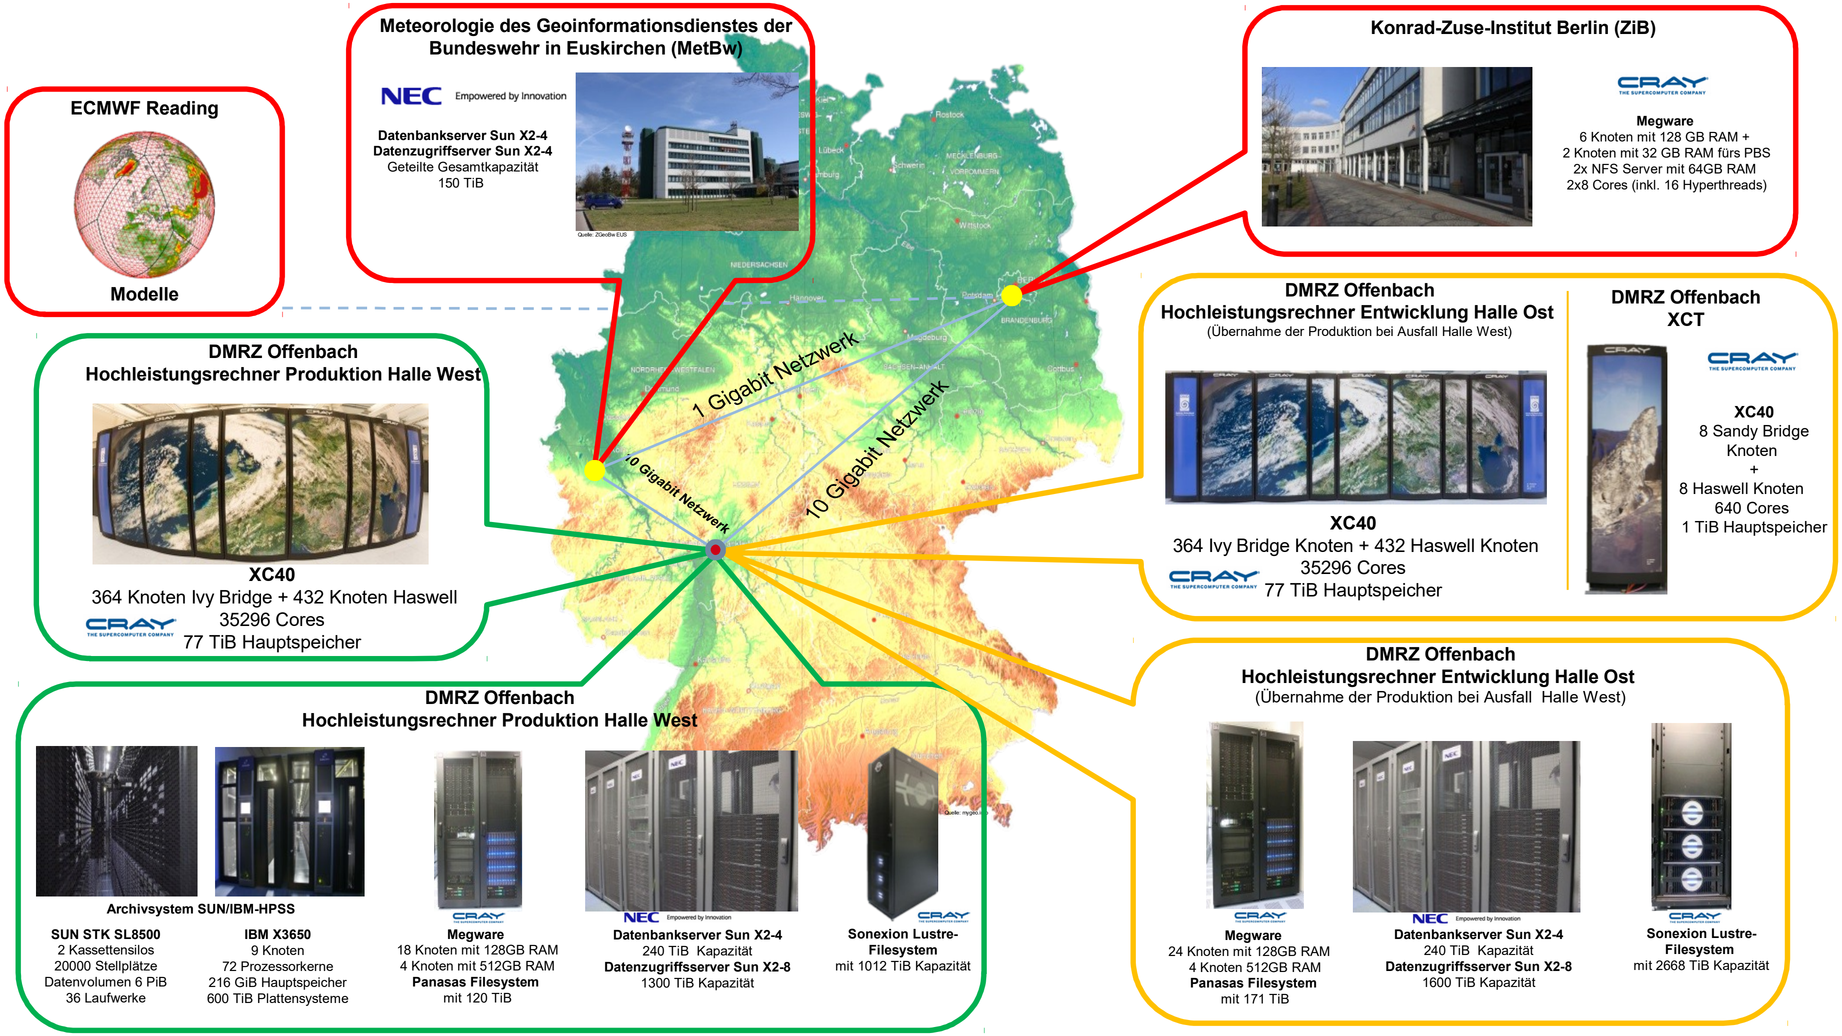

Concept: A. Pielicke, M. Jonas
Current: November 2014

Software configuration

Job management

Job dispatchment: SMS / ecFlow

- Timed job execution
- Interjob dependency
- Status reports + output capture
- Manual starts, restarts, aborts
- Transferability between halls & computing centers

Monitoring & integration

Several monitoring systems, depending on target components

- Nagios
- Icinga
- Big Brother
- Custom software

Testing and building with Jenkins

Data management overview

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

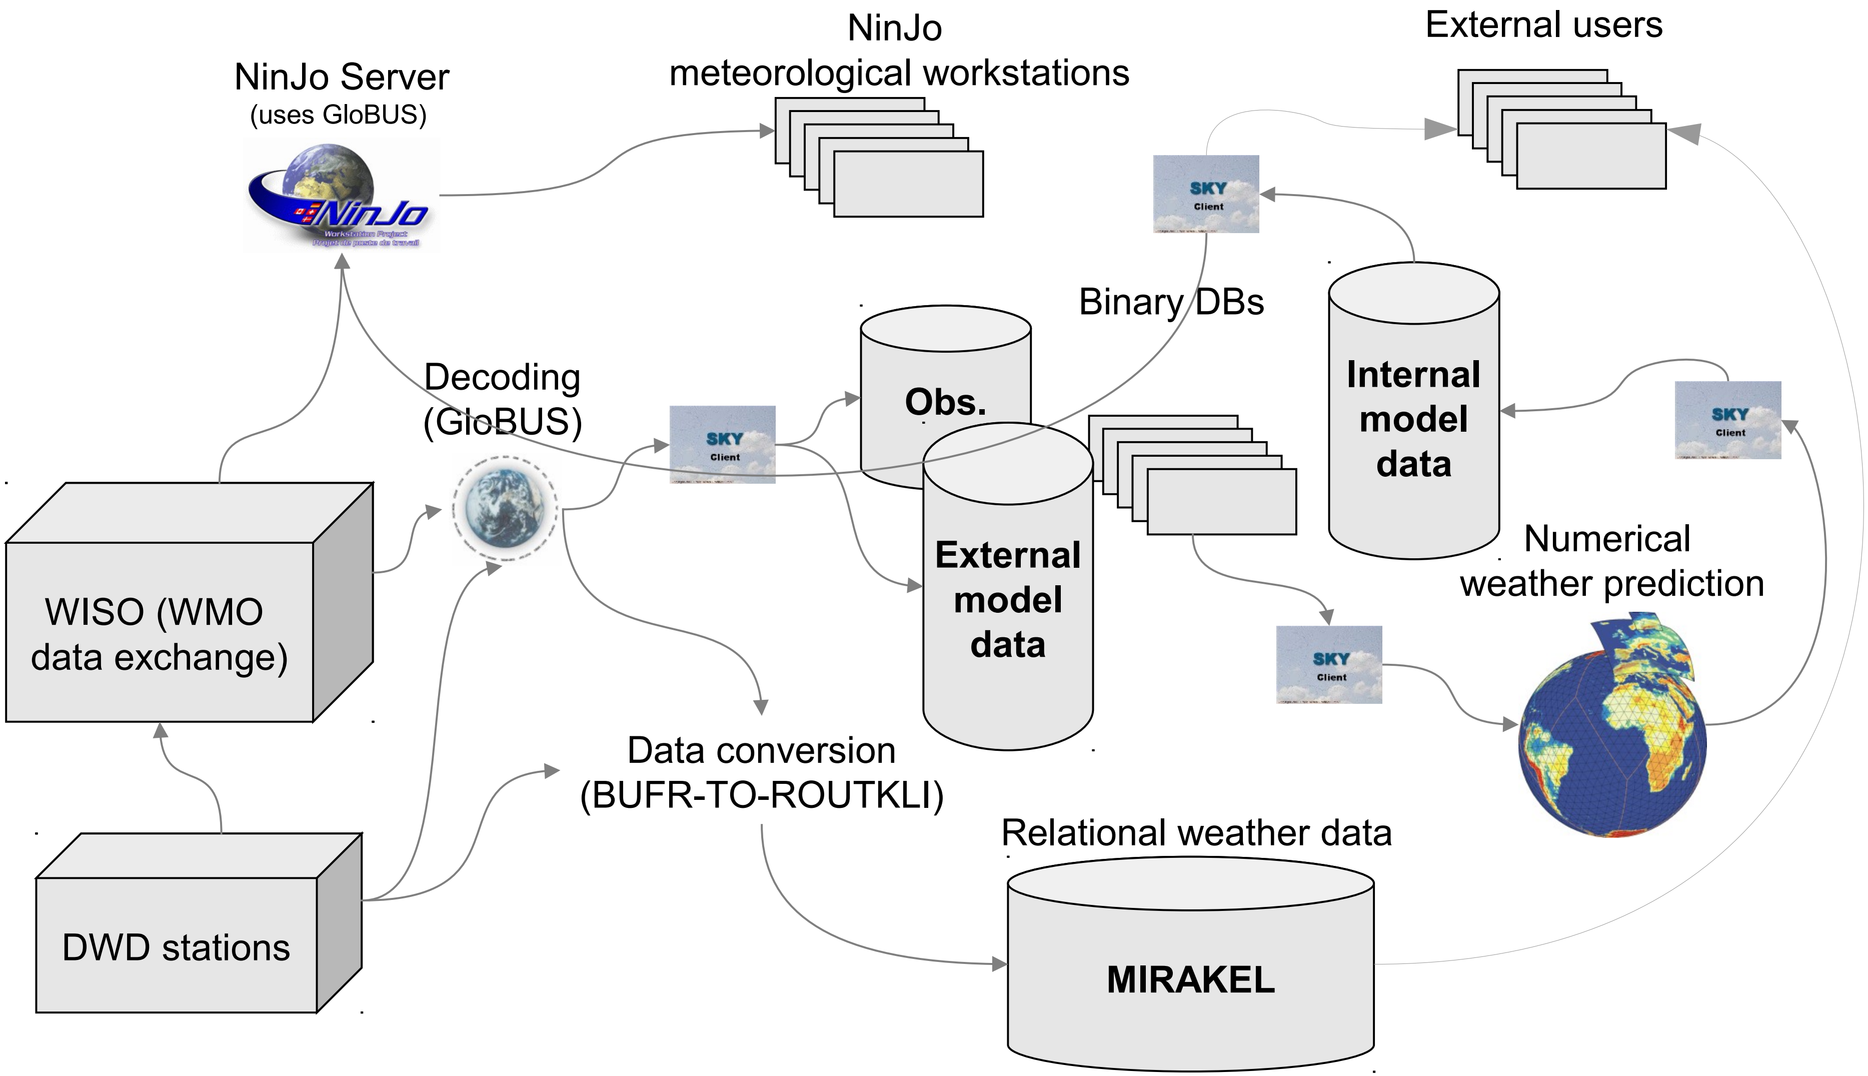

AFD: Automated file distributor

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

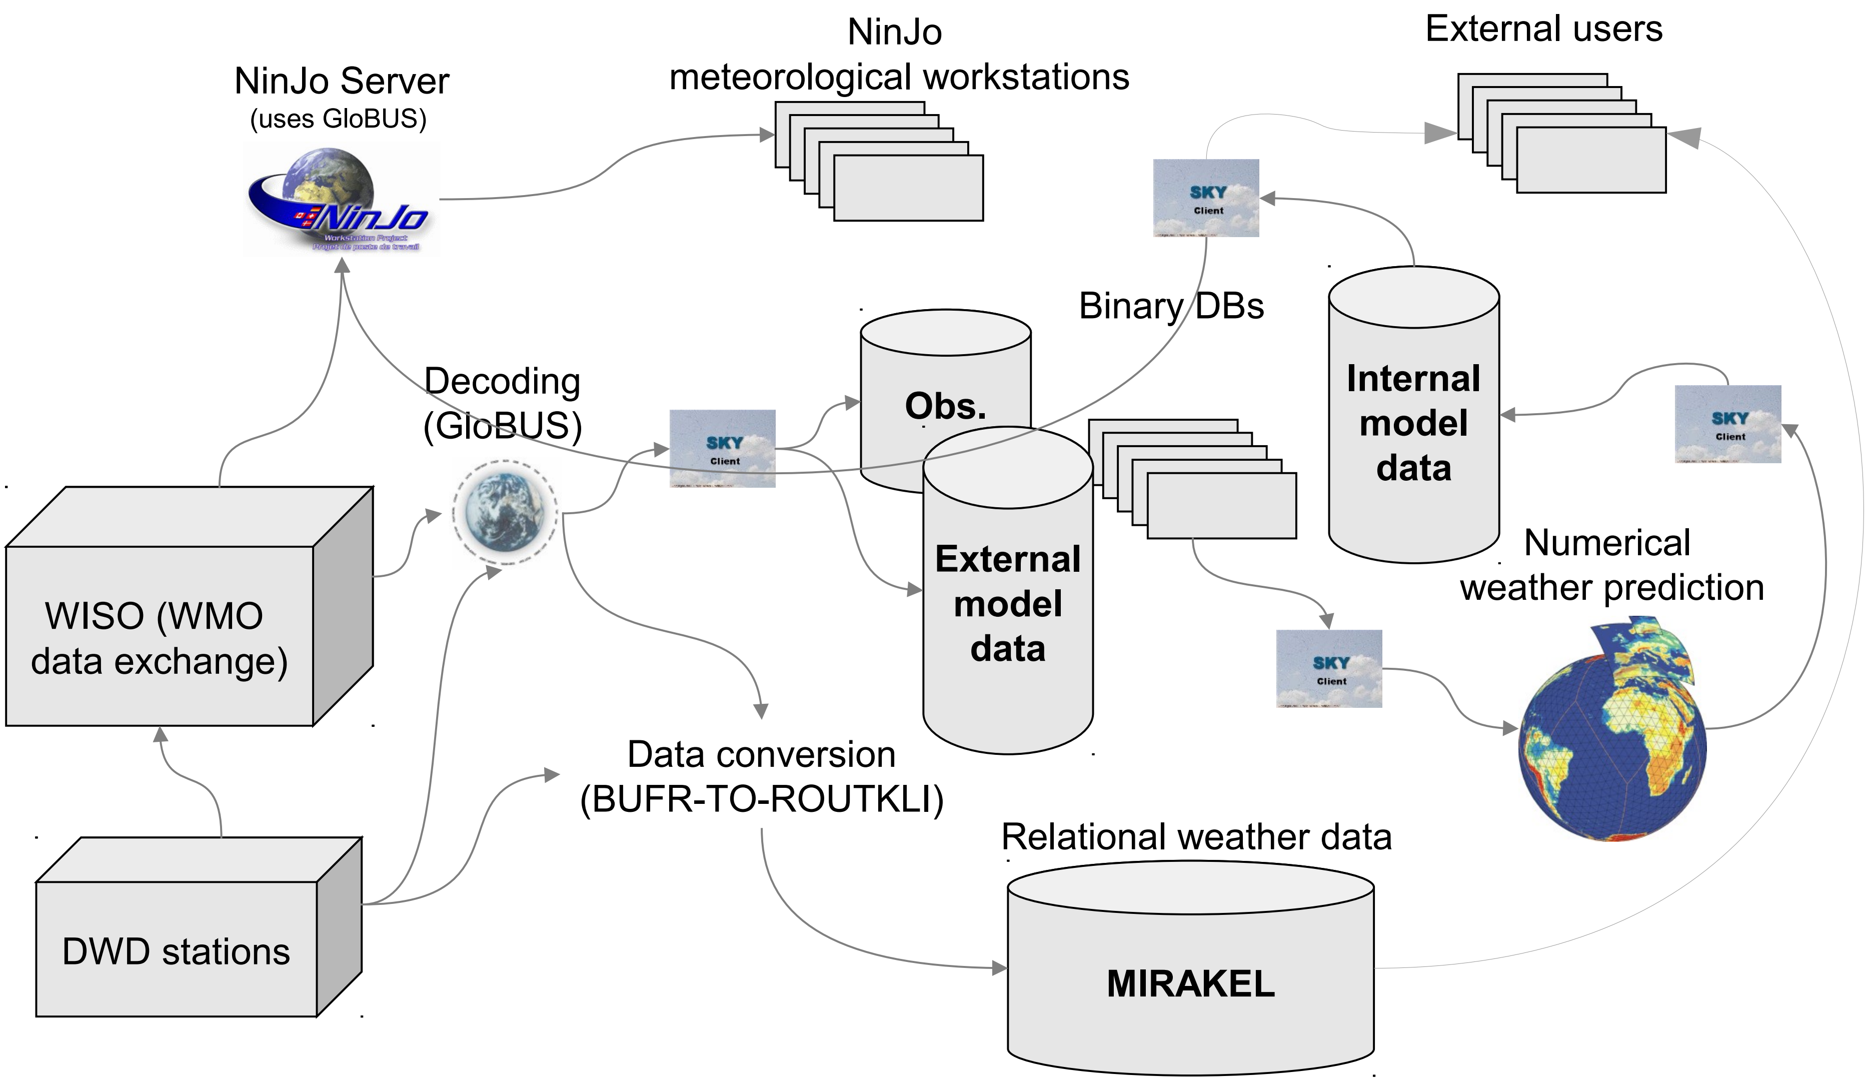

GloBUS and BUFR-TO-ROUTKLI

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

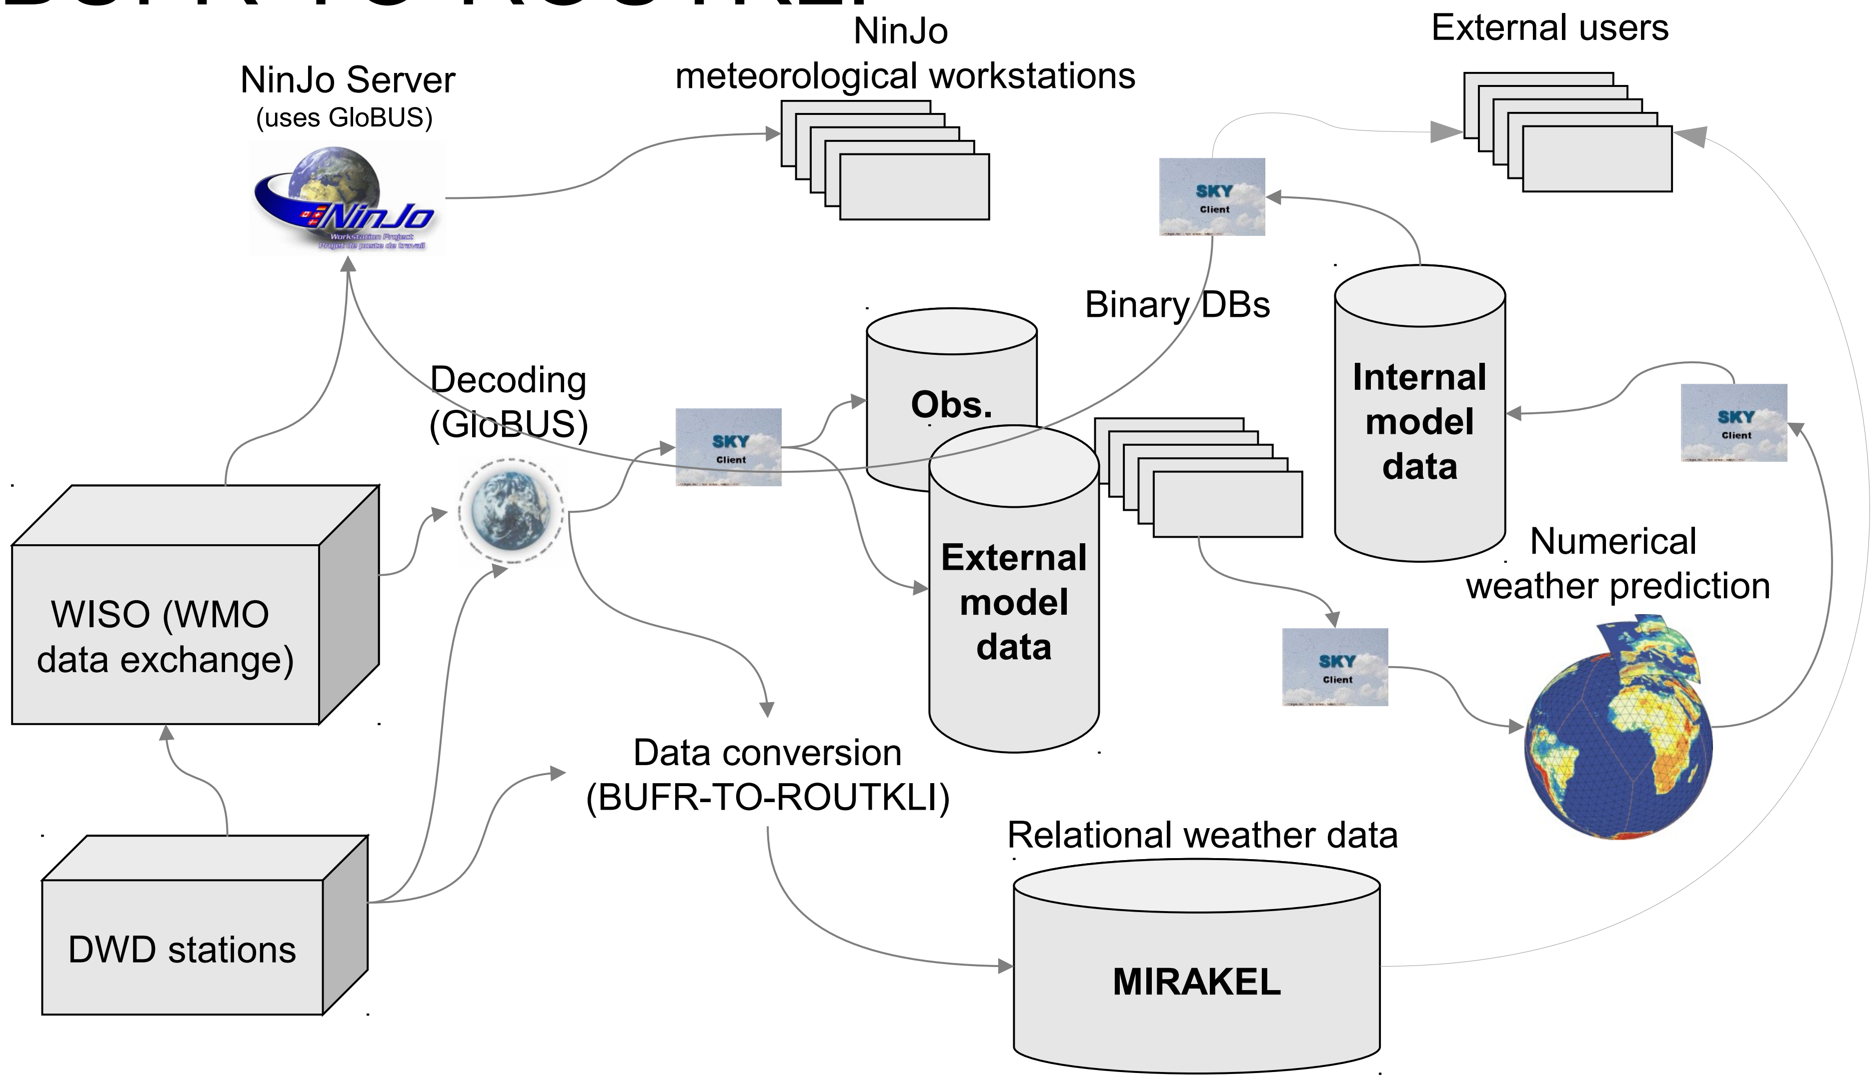

SKY

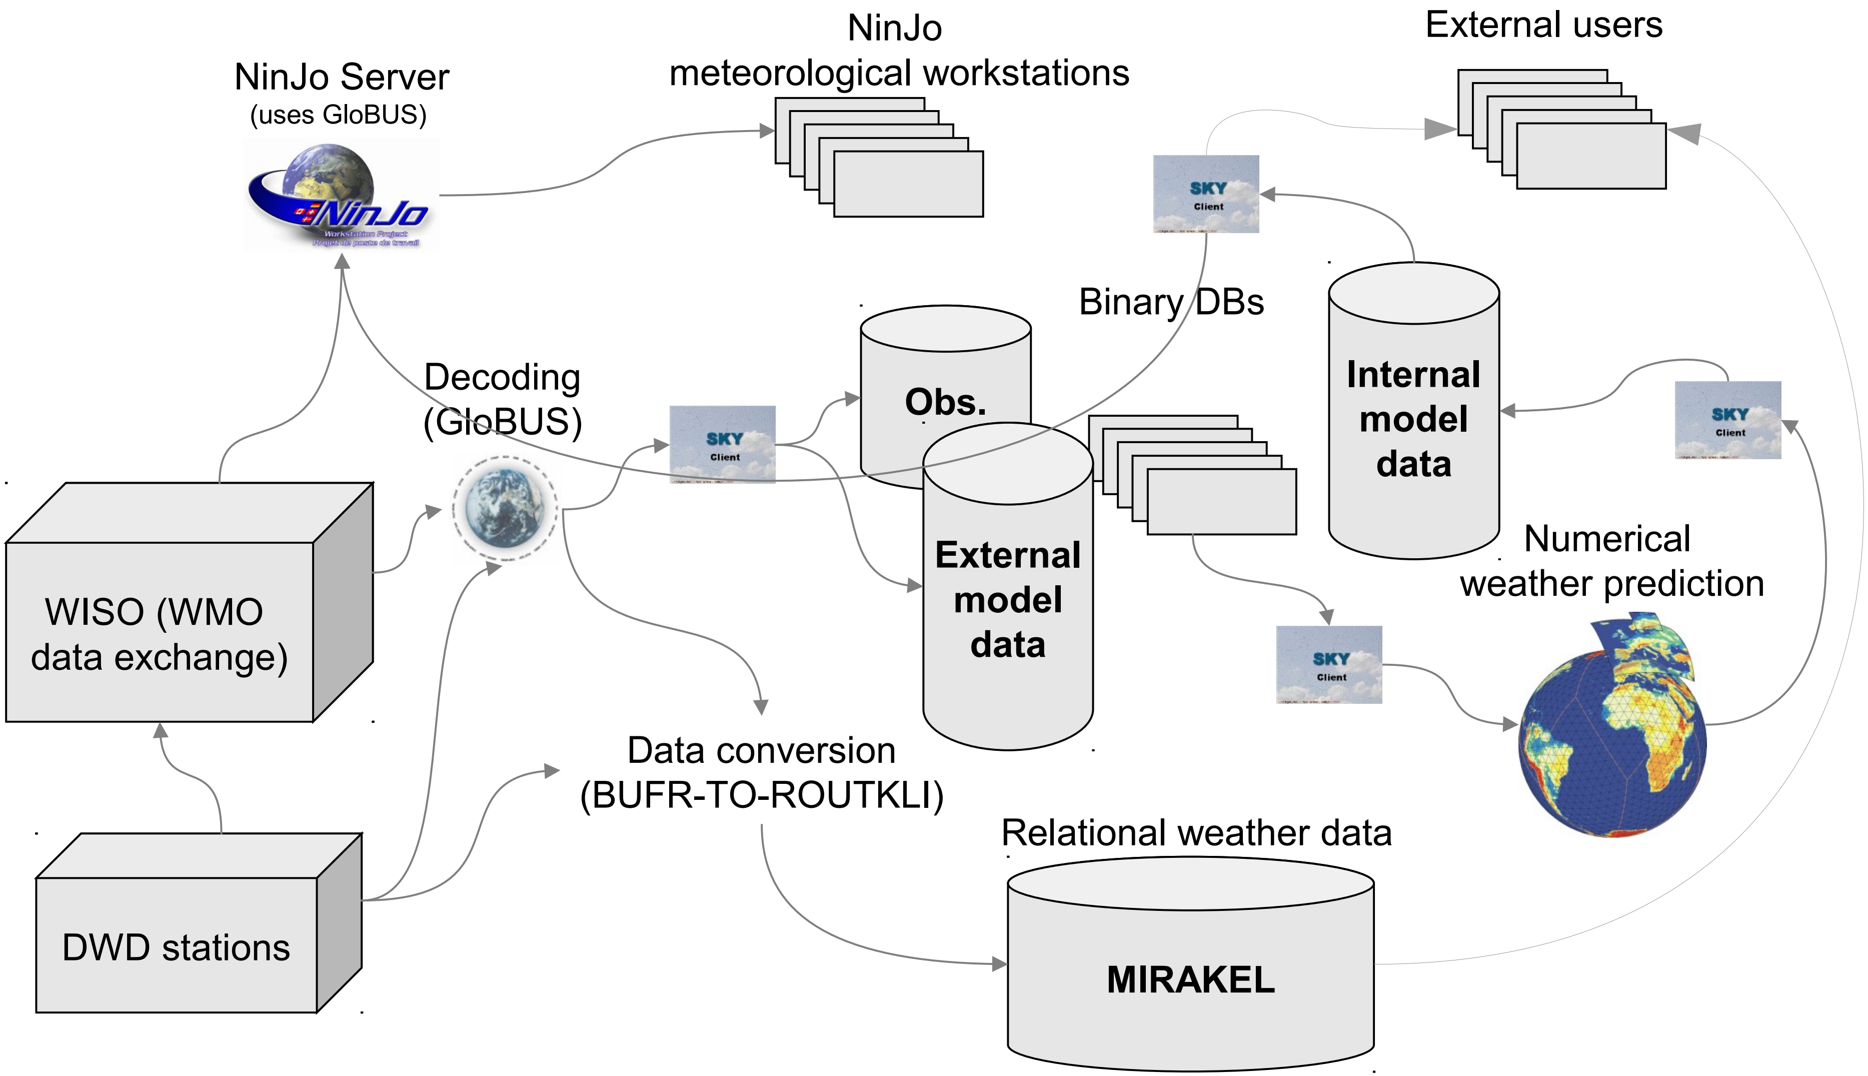
NinJo Server
(uses GloBUS)
NinJo
meteorological workstations
External users
SKY Client
Binary DBs
Decoding
(GloBUS)
Obs.
Internal
model
data
External
model
data
Numerical
weather prediction
WISO (WMO
data exchange)
Data conversion
(BUFR-TO-ROUTKLI)
Relational weather data
DWD stations
MIRAKEL

Meteorological databases

1. Time critical **(binary) databases**
Applications: NWP, postprocessing, etc.
- Near-real-time availability
- proprietary (in-house software SKY)
- Metadata in Oracle-RDBMS, payload data in files

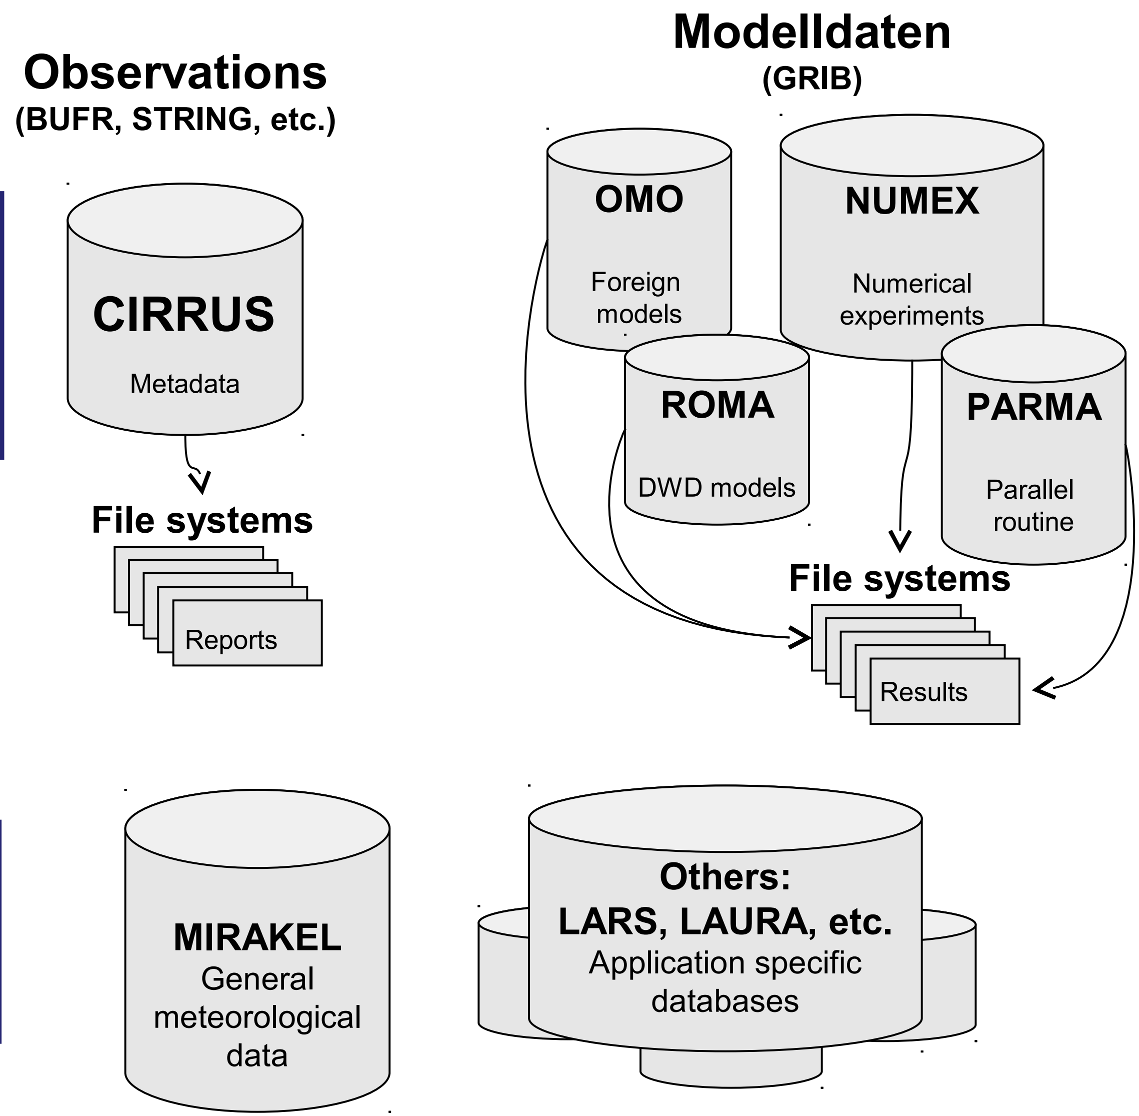

2. Relational databases
Applications: climate, web services, etc.
- Lower availability
- Easier access
- Data and metadata stored together in-DB

Access patterns

- 60,000 – 92,000 files saved per day
- 4,200 – 12,000 files accessed per day
- 35 – 75 TiB saved per day
- 5 – 17 TiB accessed per day
- 280,000 – 950,000 read requests per day (often overlapping)

Geo web services

- WCS, WFS, WMS, etc. served by Geoserver
- File-based access for rasters
- Vectors stored directly in DB

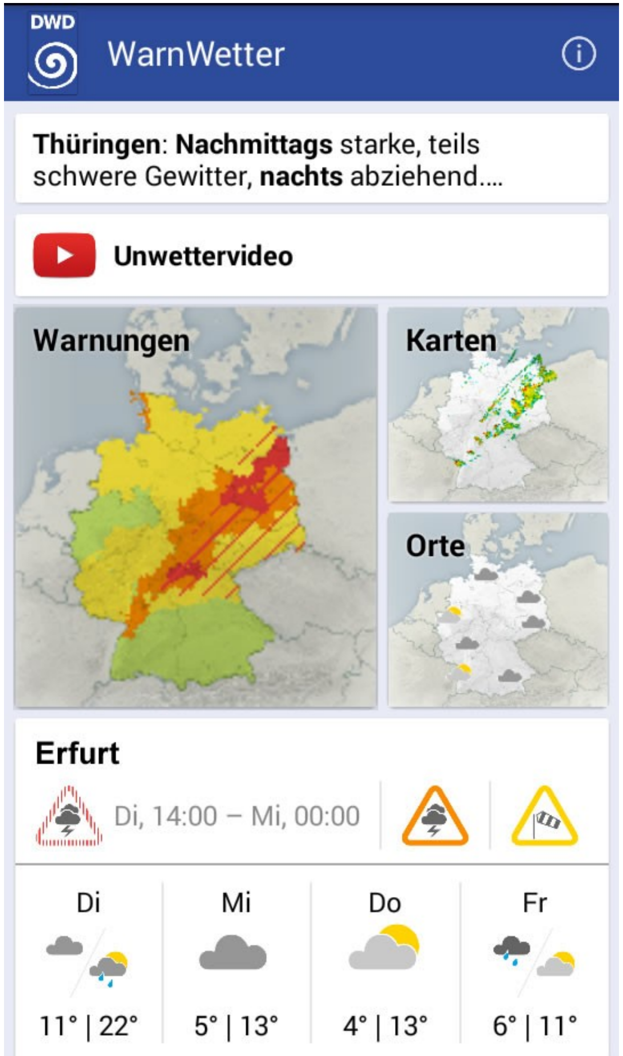

Current challenges

Observation data volume

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Example: Growth of satellite data

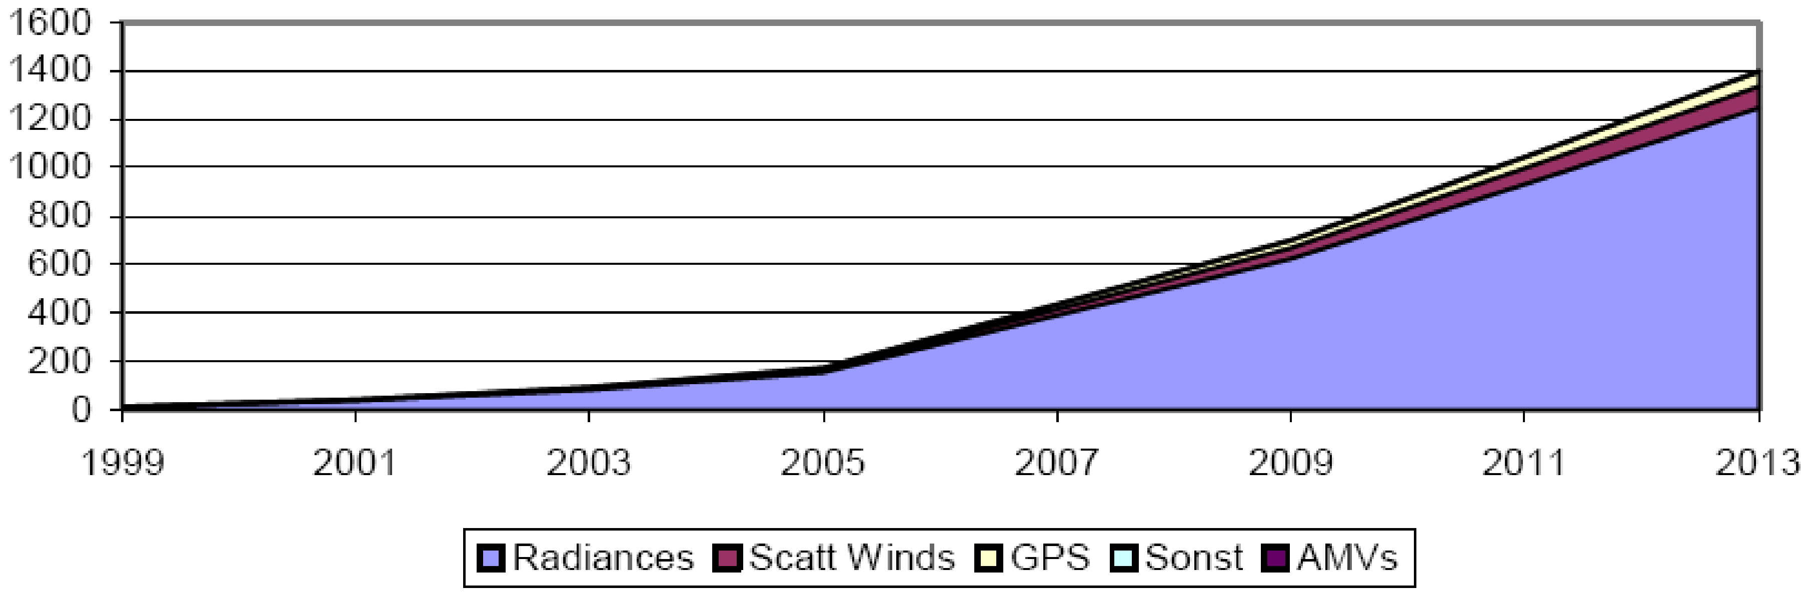

Observation data volume

Example: Growth of radiosonde data

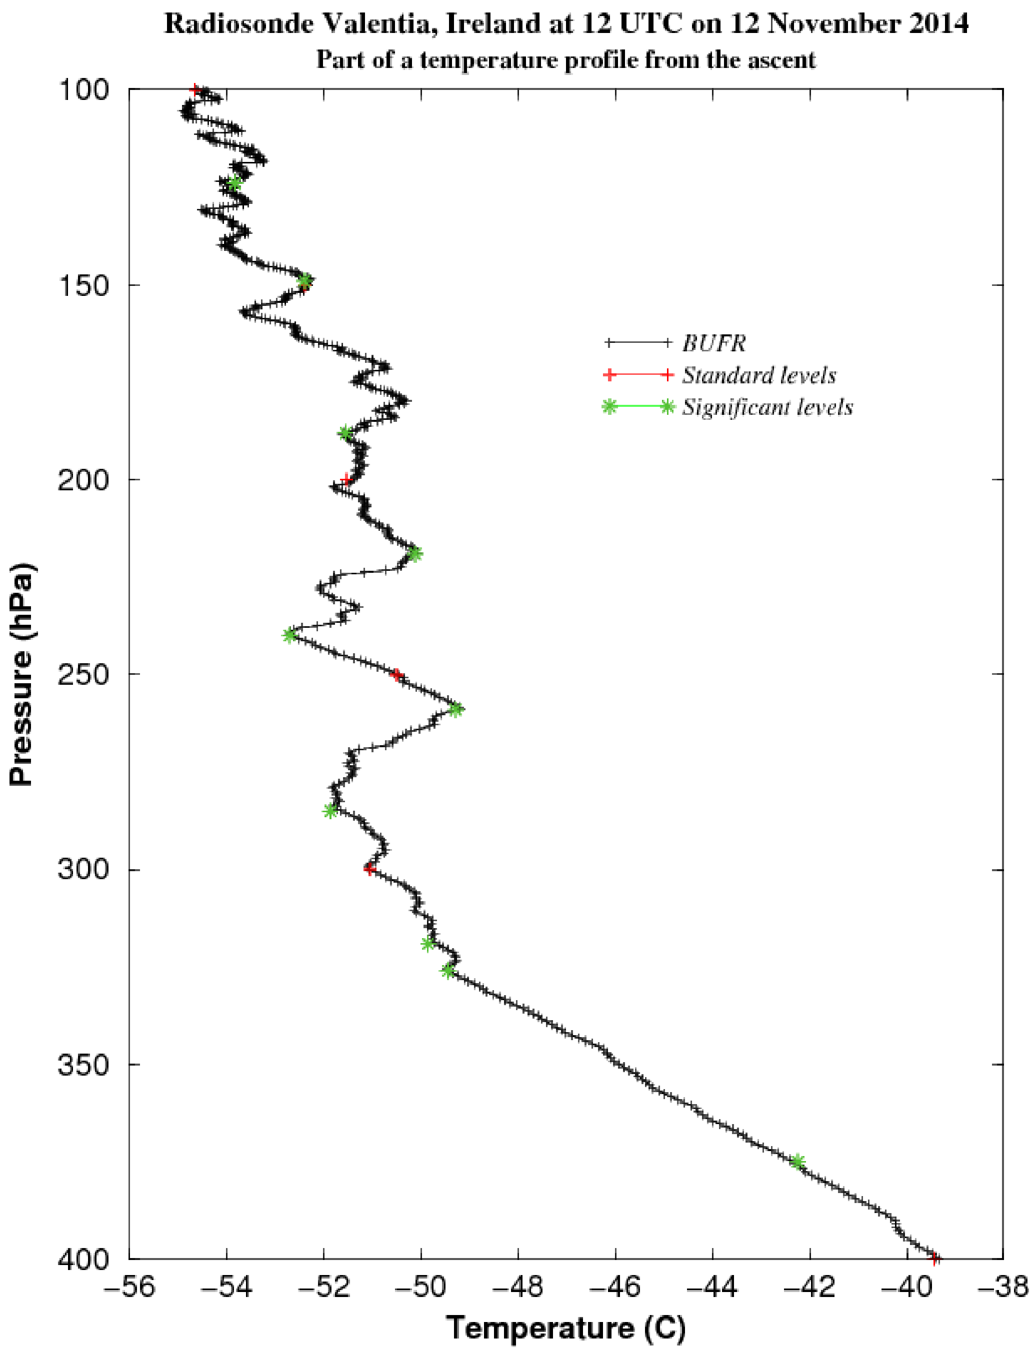

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Increase in post-processing DB access

SKY access patterns are changing:

- 2015: Factor 2.5 – 5 increase in I/O
- Some estimates for 2017 estimate increase by factor 41 (worst case scenario)

Developments in NWP

- **Tiles** (high resolution ground level modeling)
- **EDA** (ensemble data assimilation)
- **Ensemble ICON** (global & EU)
- **COSMO-D2**(-EPS, 40 mem.): 3.2x more data (higher resolution in all dims)
- **KENDA** (km-scale Ensemble Data Assimilation)
- More ensemble members

Increased contact with external users

More and more partner networks want to provide data in non-standard formats (XML, NetCDF, etc.)

Custom applications for end users also require output data in new formats (GeoTiFF, etc.)

Many users want higher data volumes or data on-demand

Future plans

HPC expansion

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Post-acquisition is pre-acquisition

Data access optimization

Primary interface to file system, SKY, has several extensions in the pipeline:

- Higher parallelization
- More intelligent caching strategies

Encoding / decoding are also being reworked to move away from monolithic to more disparate software components

Separation of concerns

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

New projects serving data to end users will be increasingly quarantined from production resources

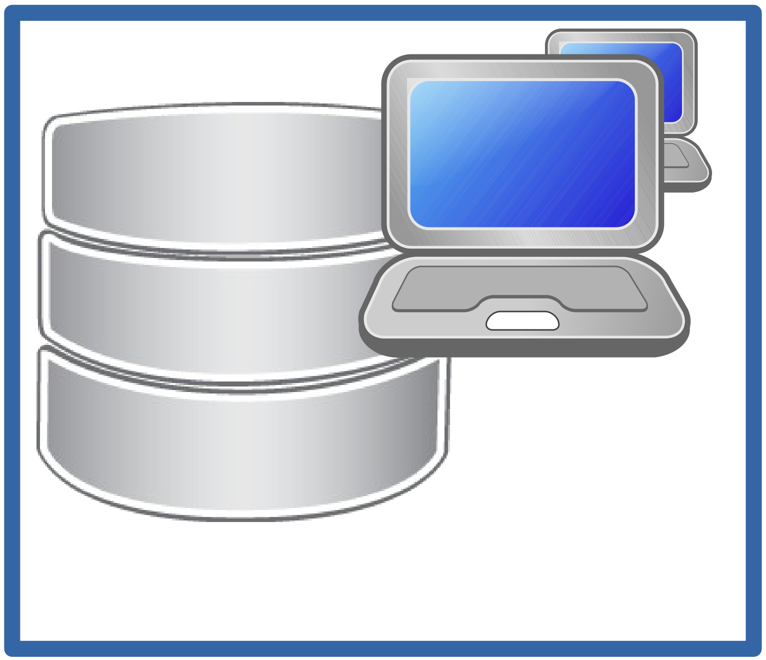

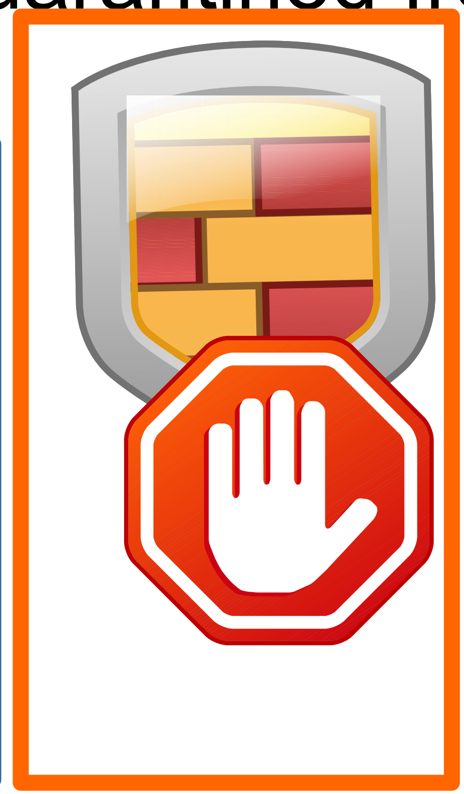

Cooperation and coordination: IT and research

Thanks!

Questions?