




# IO-SEA Hackathon

Hands-On Training with Data Access & Storage Interface (DASI)



**EuroHPC**  
Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955811. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, the Czech Republic, Germany, Ireland, Sweden, and the United Kingdom.

# IO-SEA Hackathon

Hands-On Training with Data Access & Storage Interface (DASI)



**13:30** Introduction:

- ECMWF and the importance of semantic data management
- IOSEA and DASI

**14:15** Coffee Break

**14:30** Hands on: archiving data with DASI

**15:15** Coffee Break

**15:30** Hands on: retrieving data with DASI

**16:00** Q&A

**James Hawkes**  
**Metin Cakircali**  
**Jenny Wong**



**EuroHPC**  
Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955811. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, the Czech Republic, Germany, Ireland, Sweden, and the United Kingdom.

# About ECMWF

## Established in 1975, Intergovernmental Organisation

- 23 Member States | 12 Cooperating States
- 350+ staff

## 24/7 operational service

- Operational NWP – 4x HRES+ENS forecasts / day
- Supporting NWS (coupled models) and businesses

## Research institution

- Experiments to continuously improve our models
- Reforecasts and Climate Reanalysis

## Operate 2 EU Copernicus Services

- Climate Change Service (C3S)
- Atmosphere Monitoring Service (CAMS)
- Support Copernicus Emergency Management Service CEMS

## Destination Earth

- Operates two Digital Twins
- Operates the DestinE Digital Twin Engine (DTE)



*Reading,  
GB*



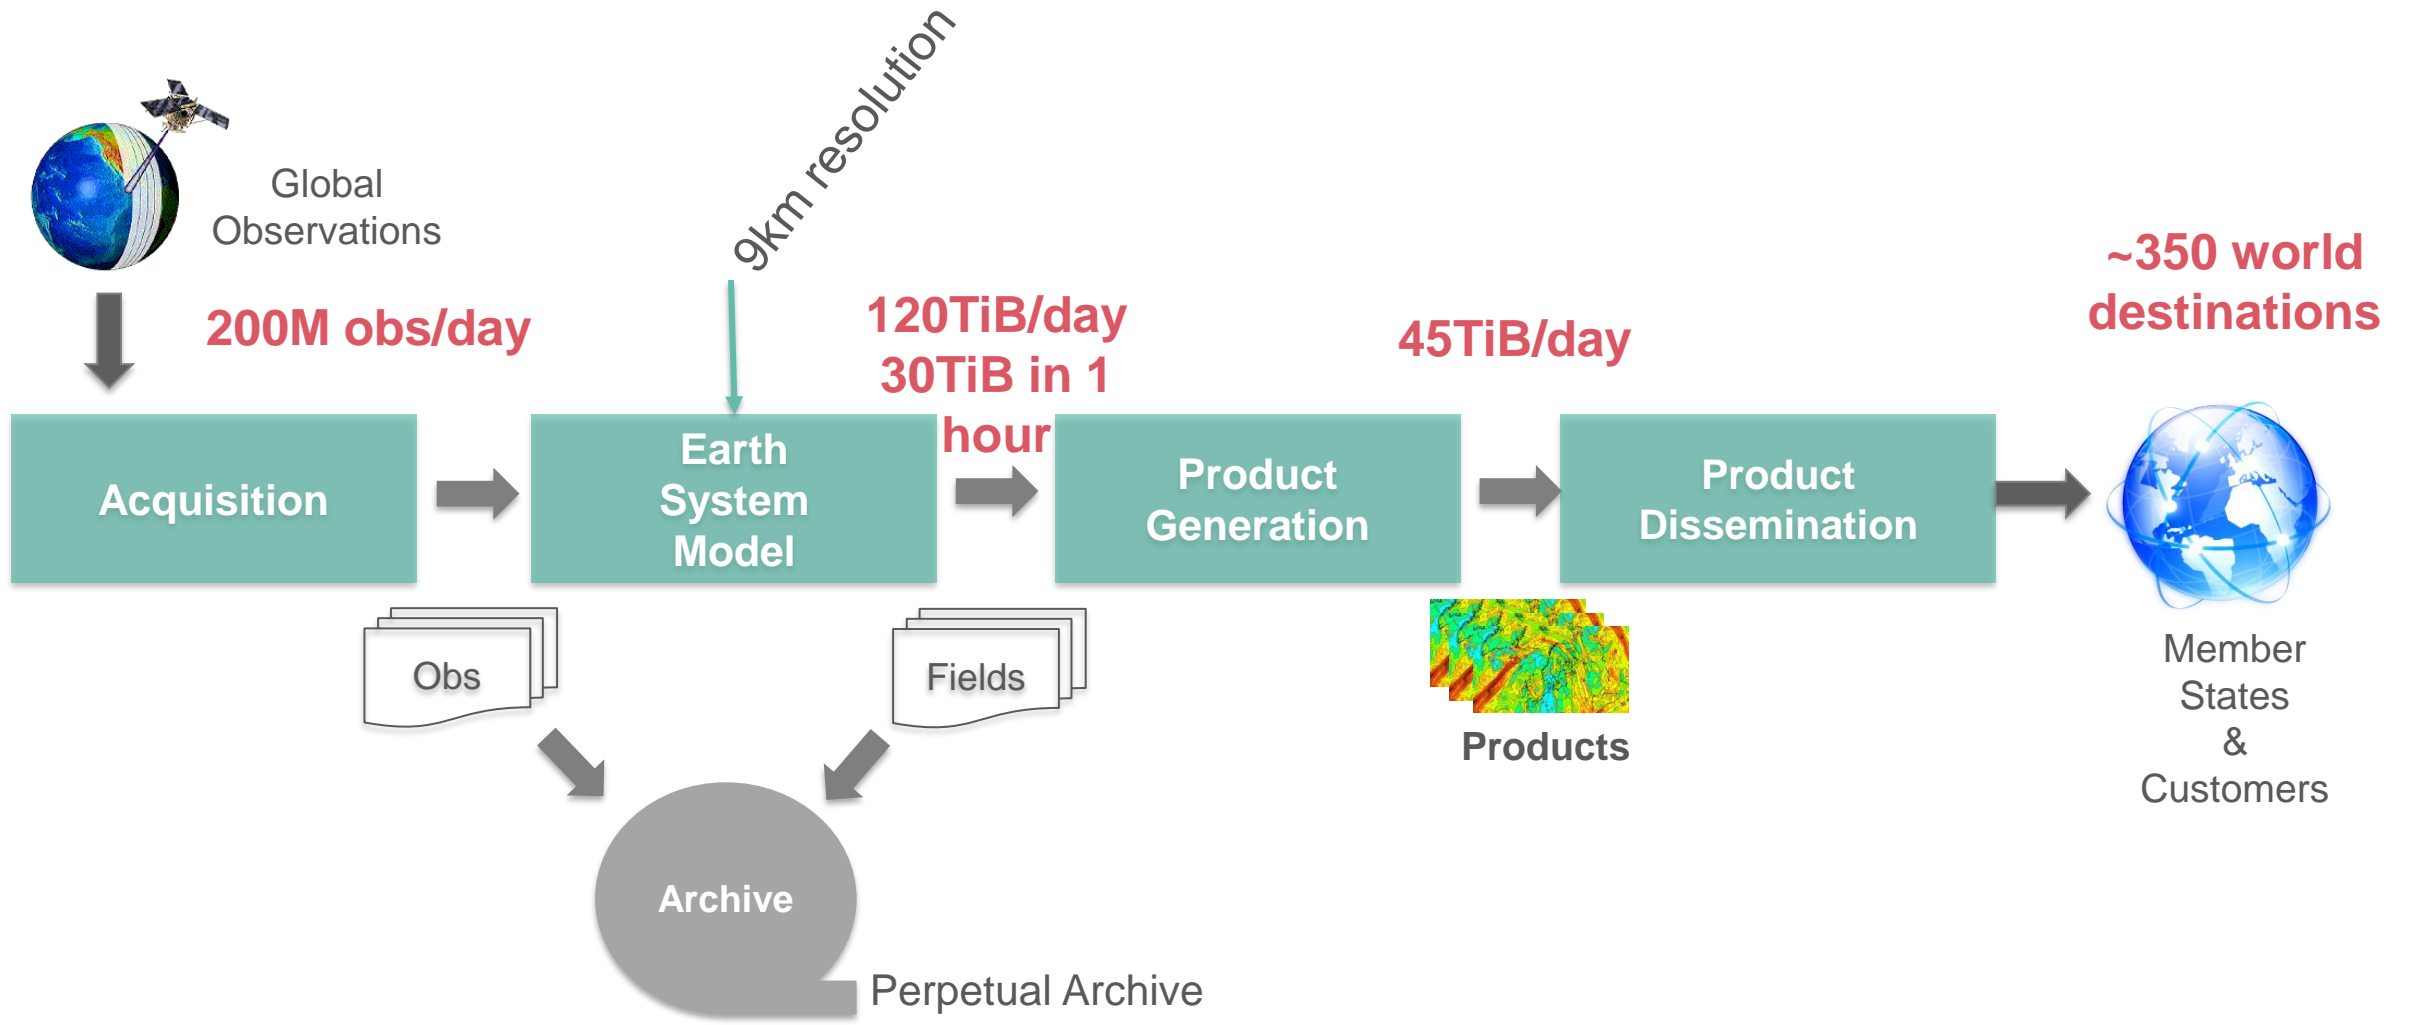
*Bonn,  
DE*



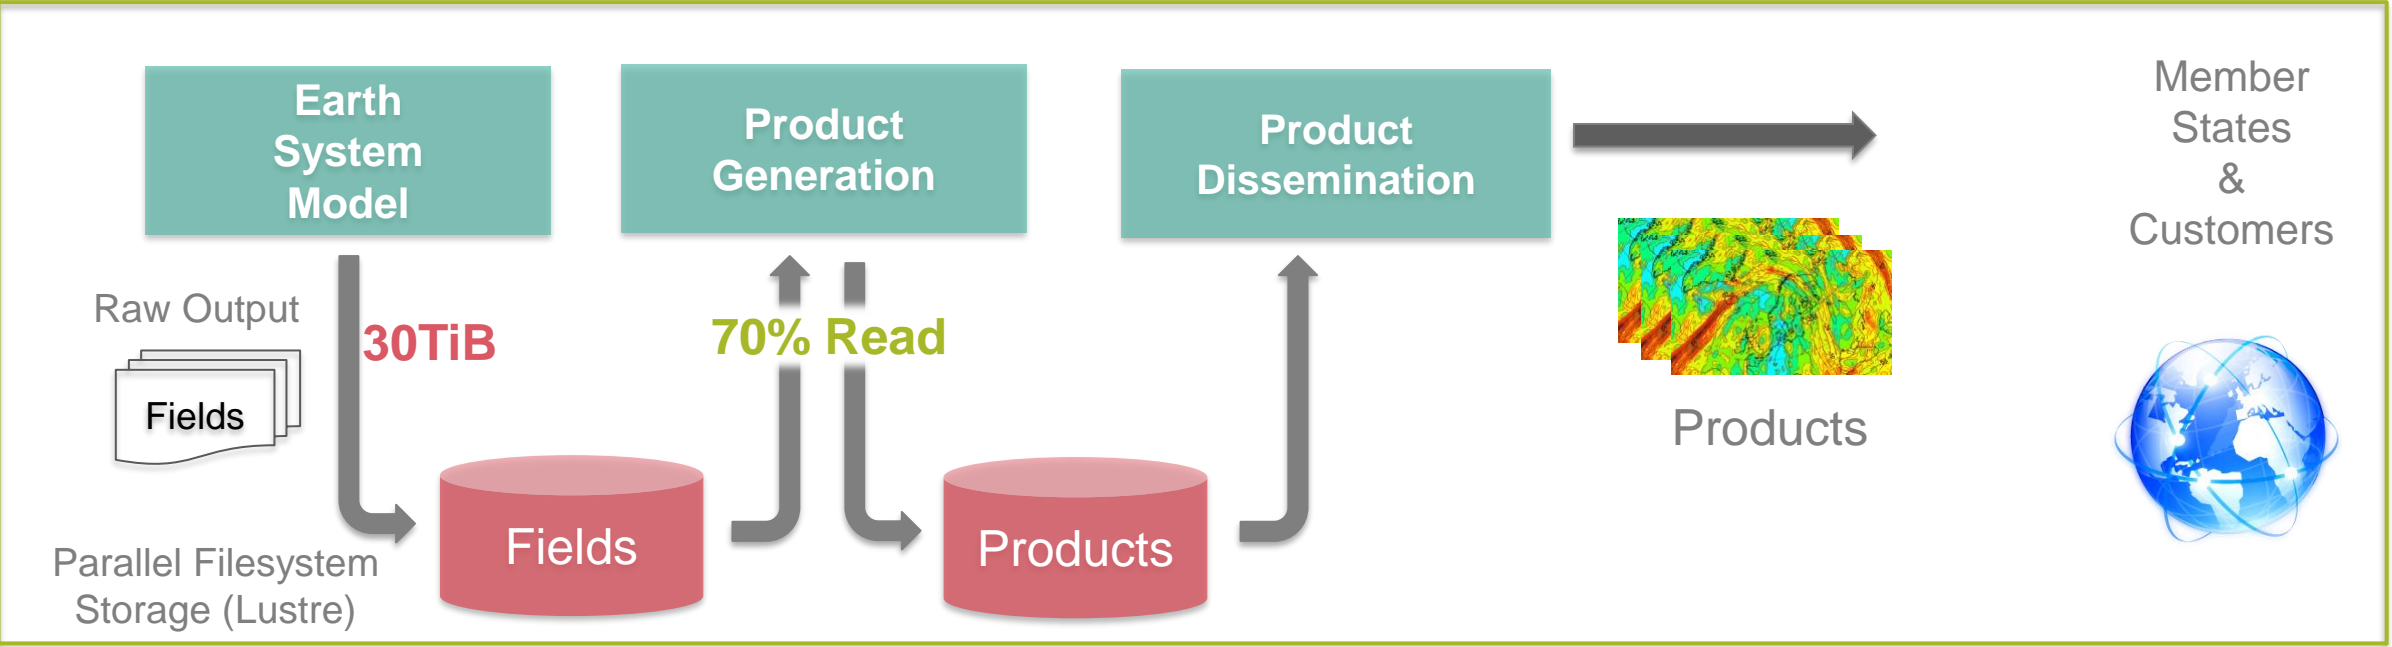
*Bologna,  
IT*



# ECMWF's Production Workflow




# ECMWF's Production Workflow



Time critical path = 1 hour window

# ECMWF's Production Workflow



# Semantic Data Management

- Data is indexed by its scientific metadata, according to a hierarchical schema
- The key used to index data carries scientific meaning
  - Not just a UUID
  - Not just storing metadata with data
  - The metadata is **used to index and uniquely identify the data**
- ECMWF archive from 1984-2023 (>400PiB) is all addressed with the same data language

## Non-semantic key

8s09sno5tdyjopj92asy23

## Semantic key

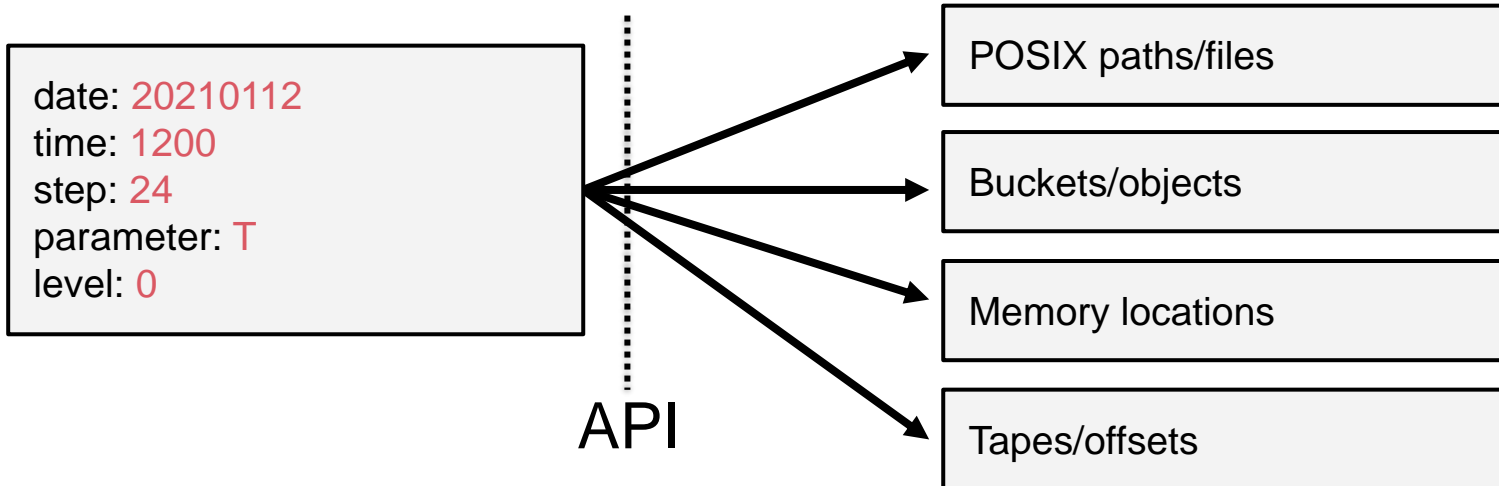
date: 20210112  
time: 1200  
step: 24  
parameter: T  
level: 0

# Semantic Data Management

- The most basic semantic data access can be done with files...

```
./../20210112/1200/24/0/T/...
```

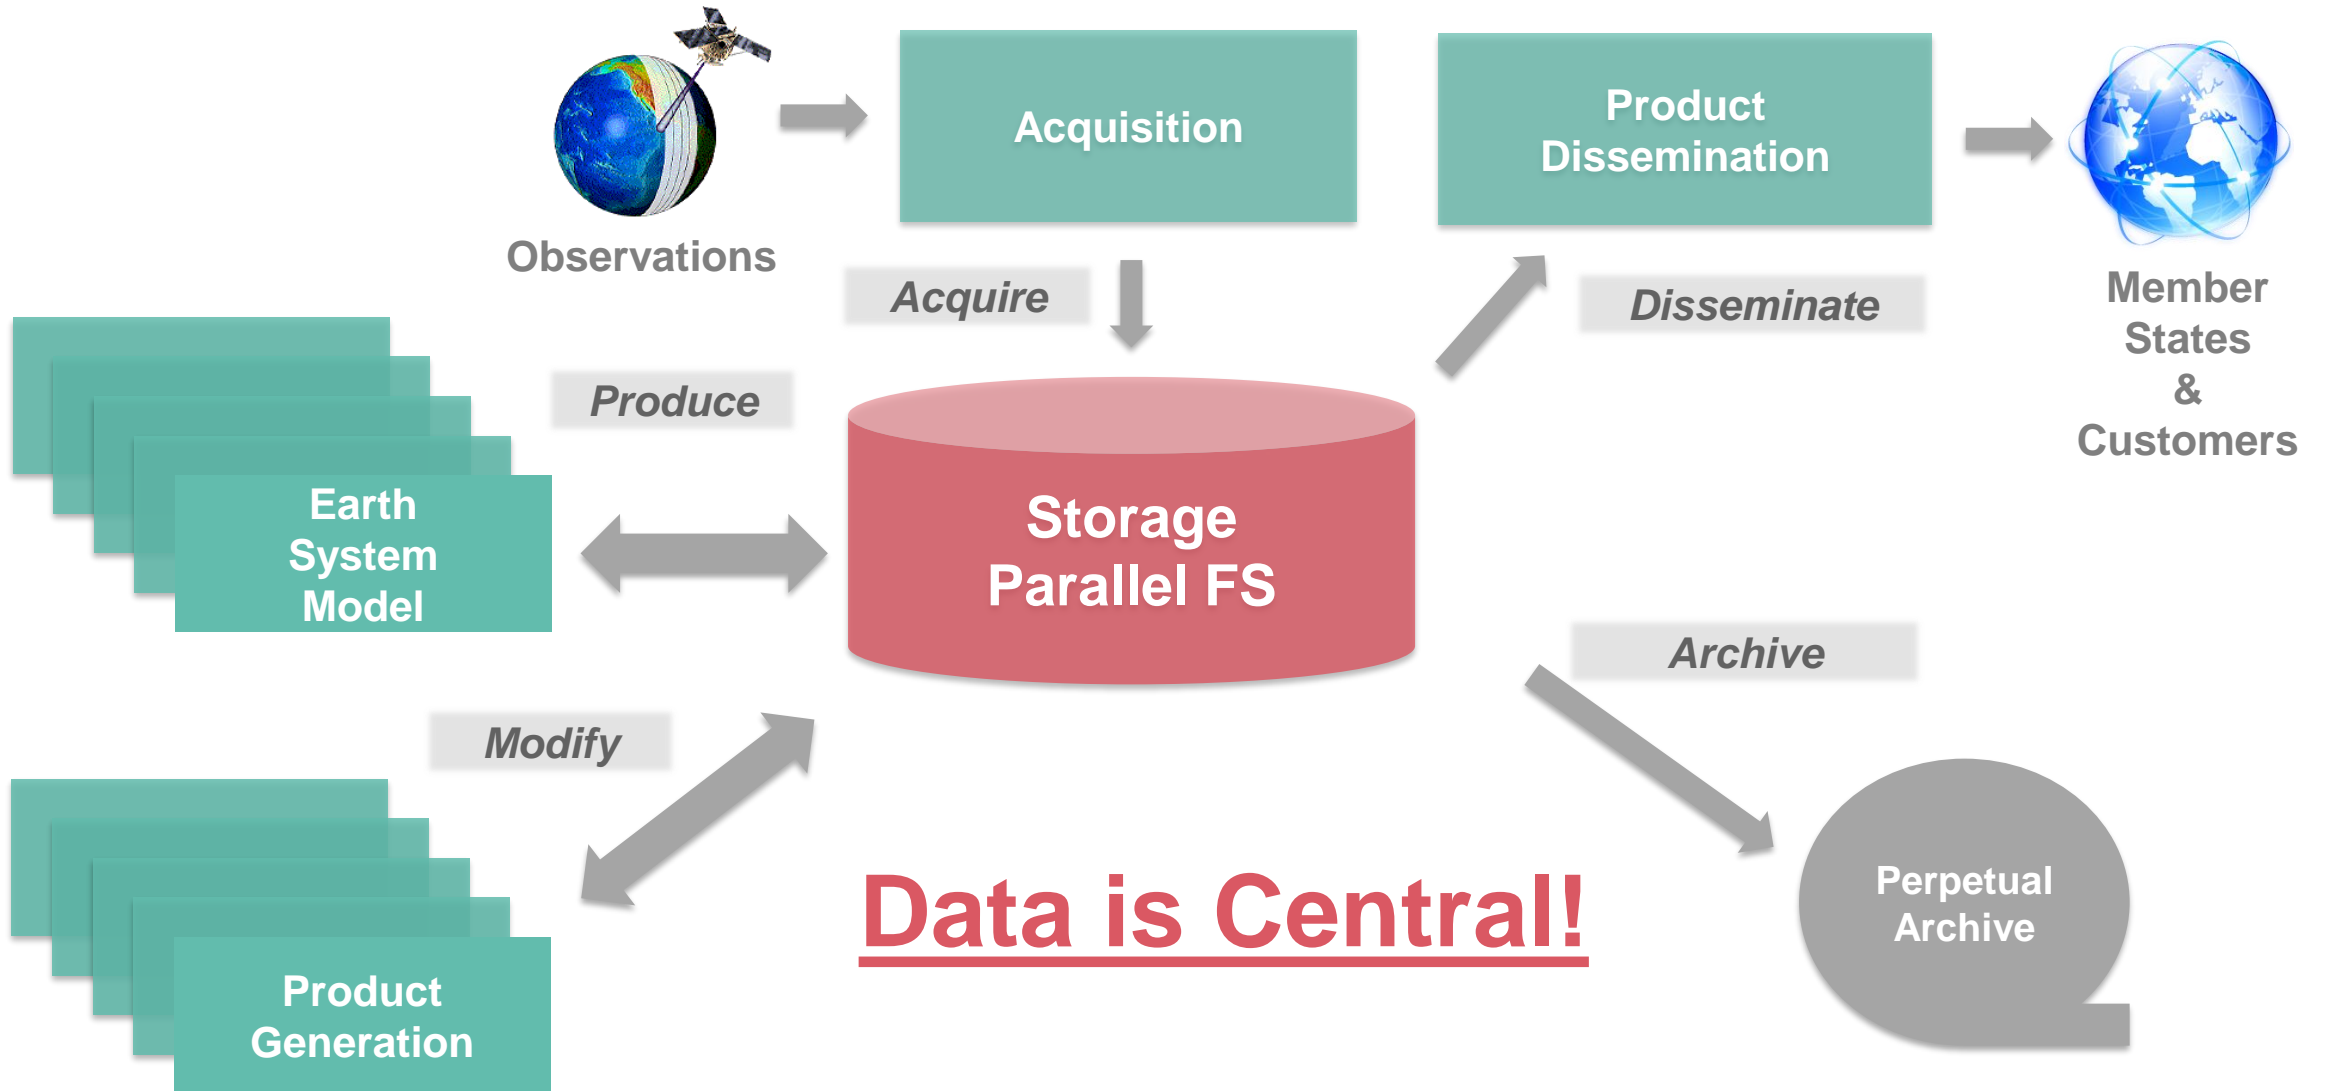
- ... but a proper implementation decouples the scientific identification from the storage resource



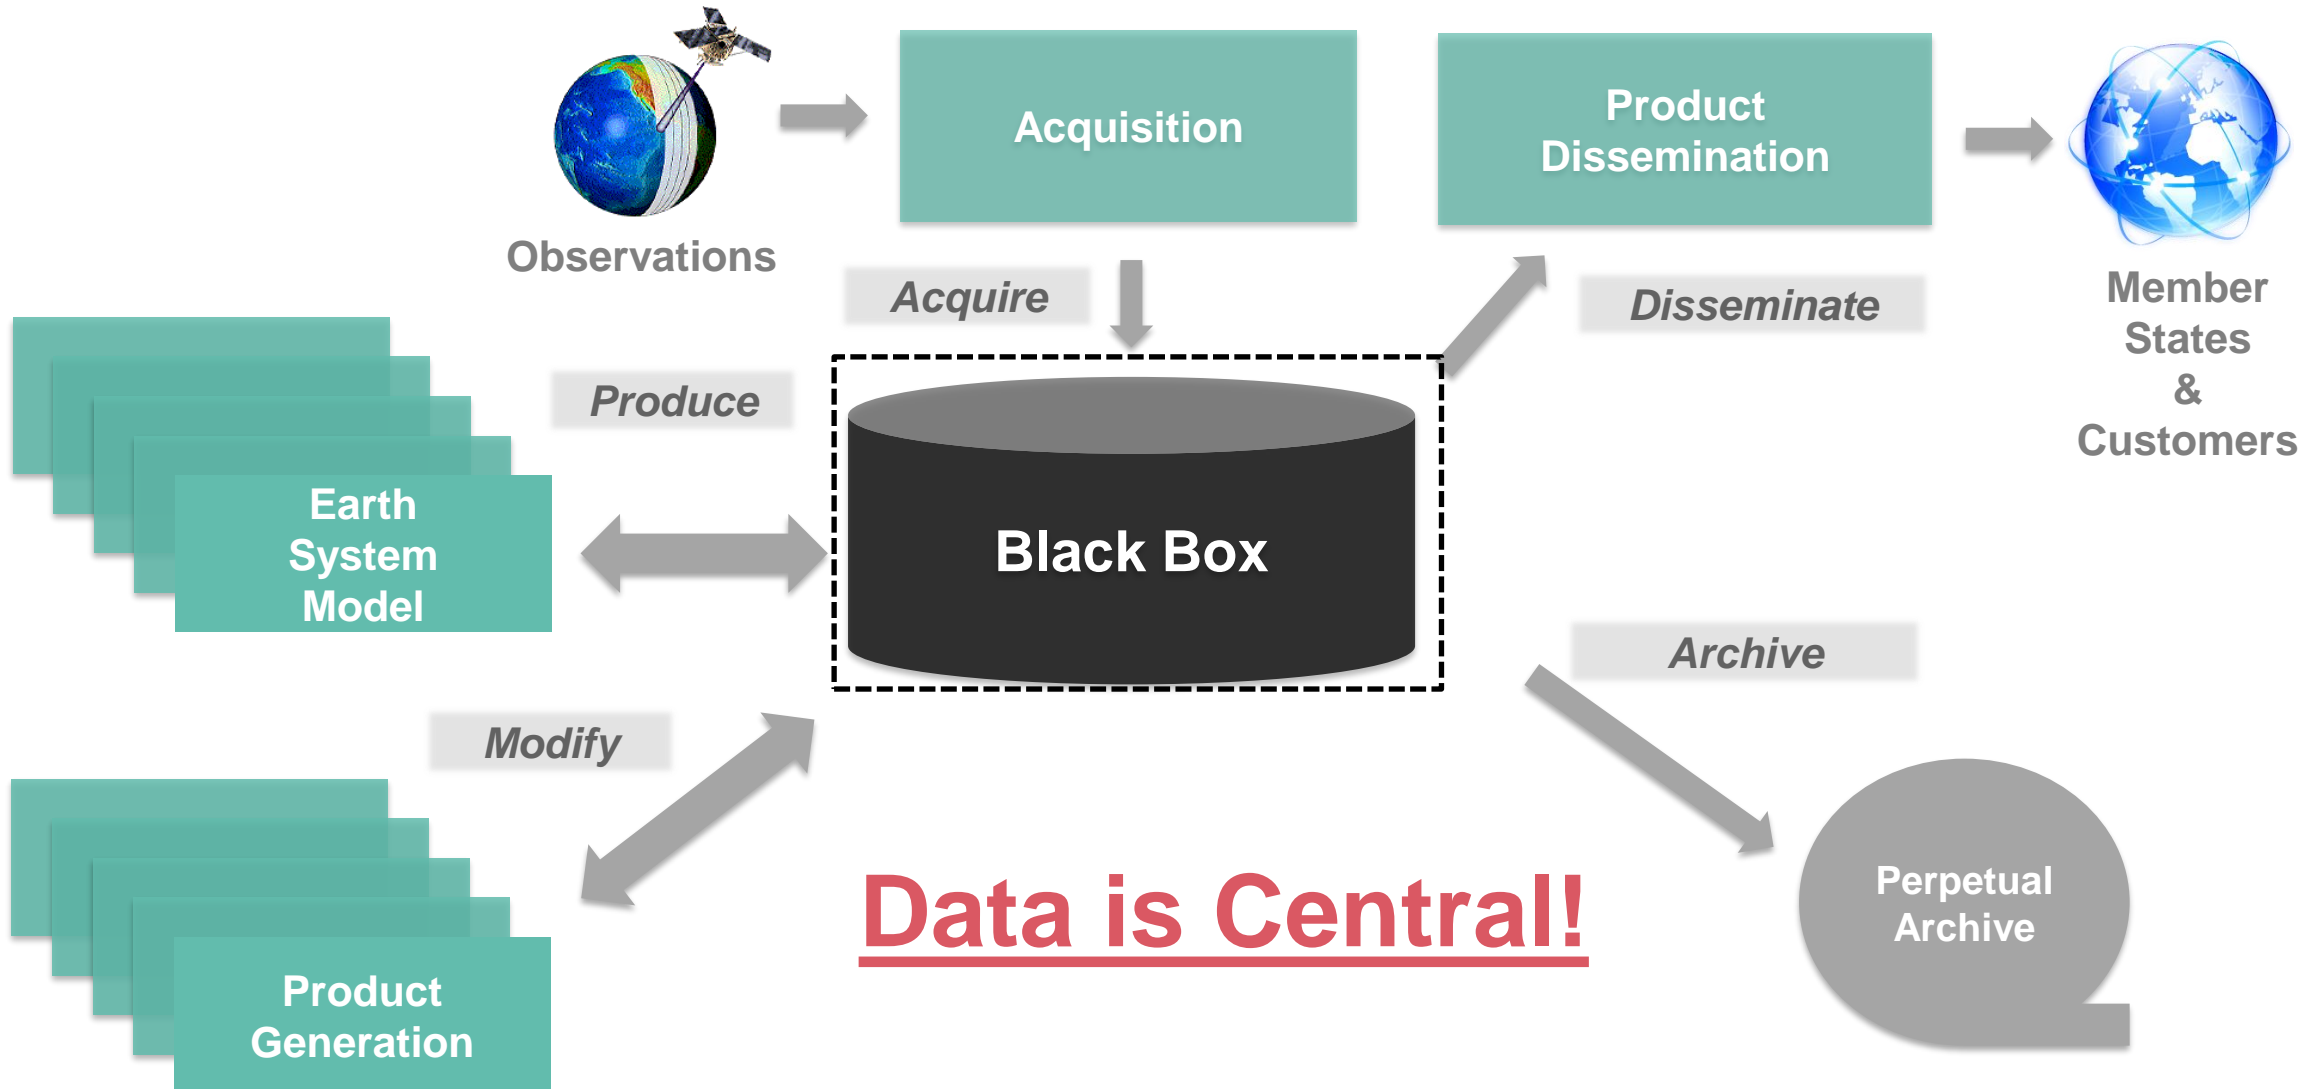
- ... and the applications don't need to care how the objects are stored.



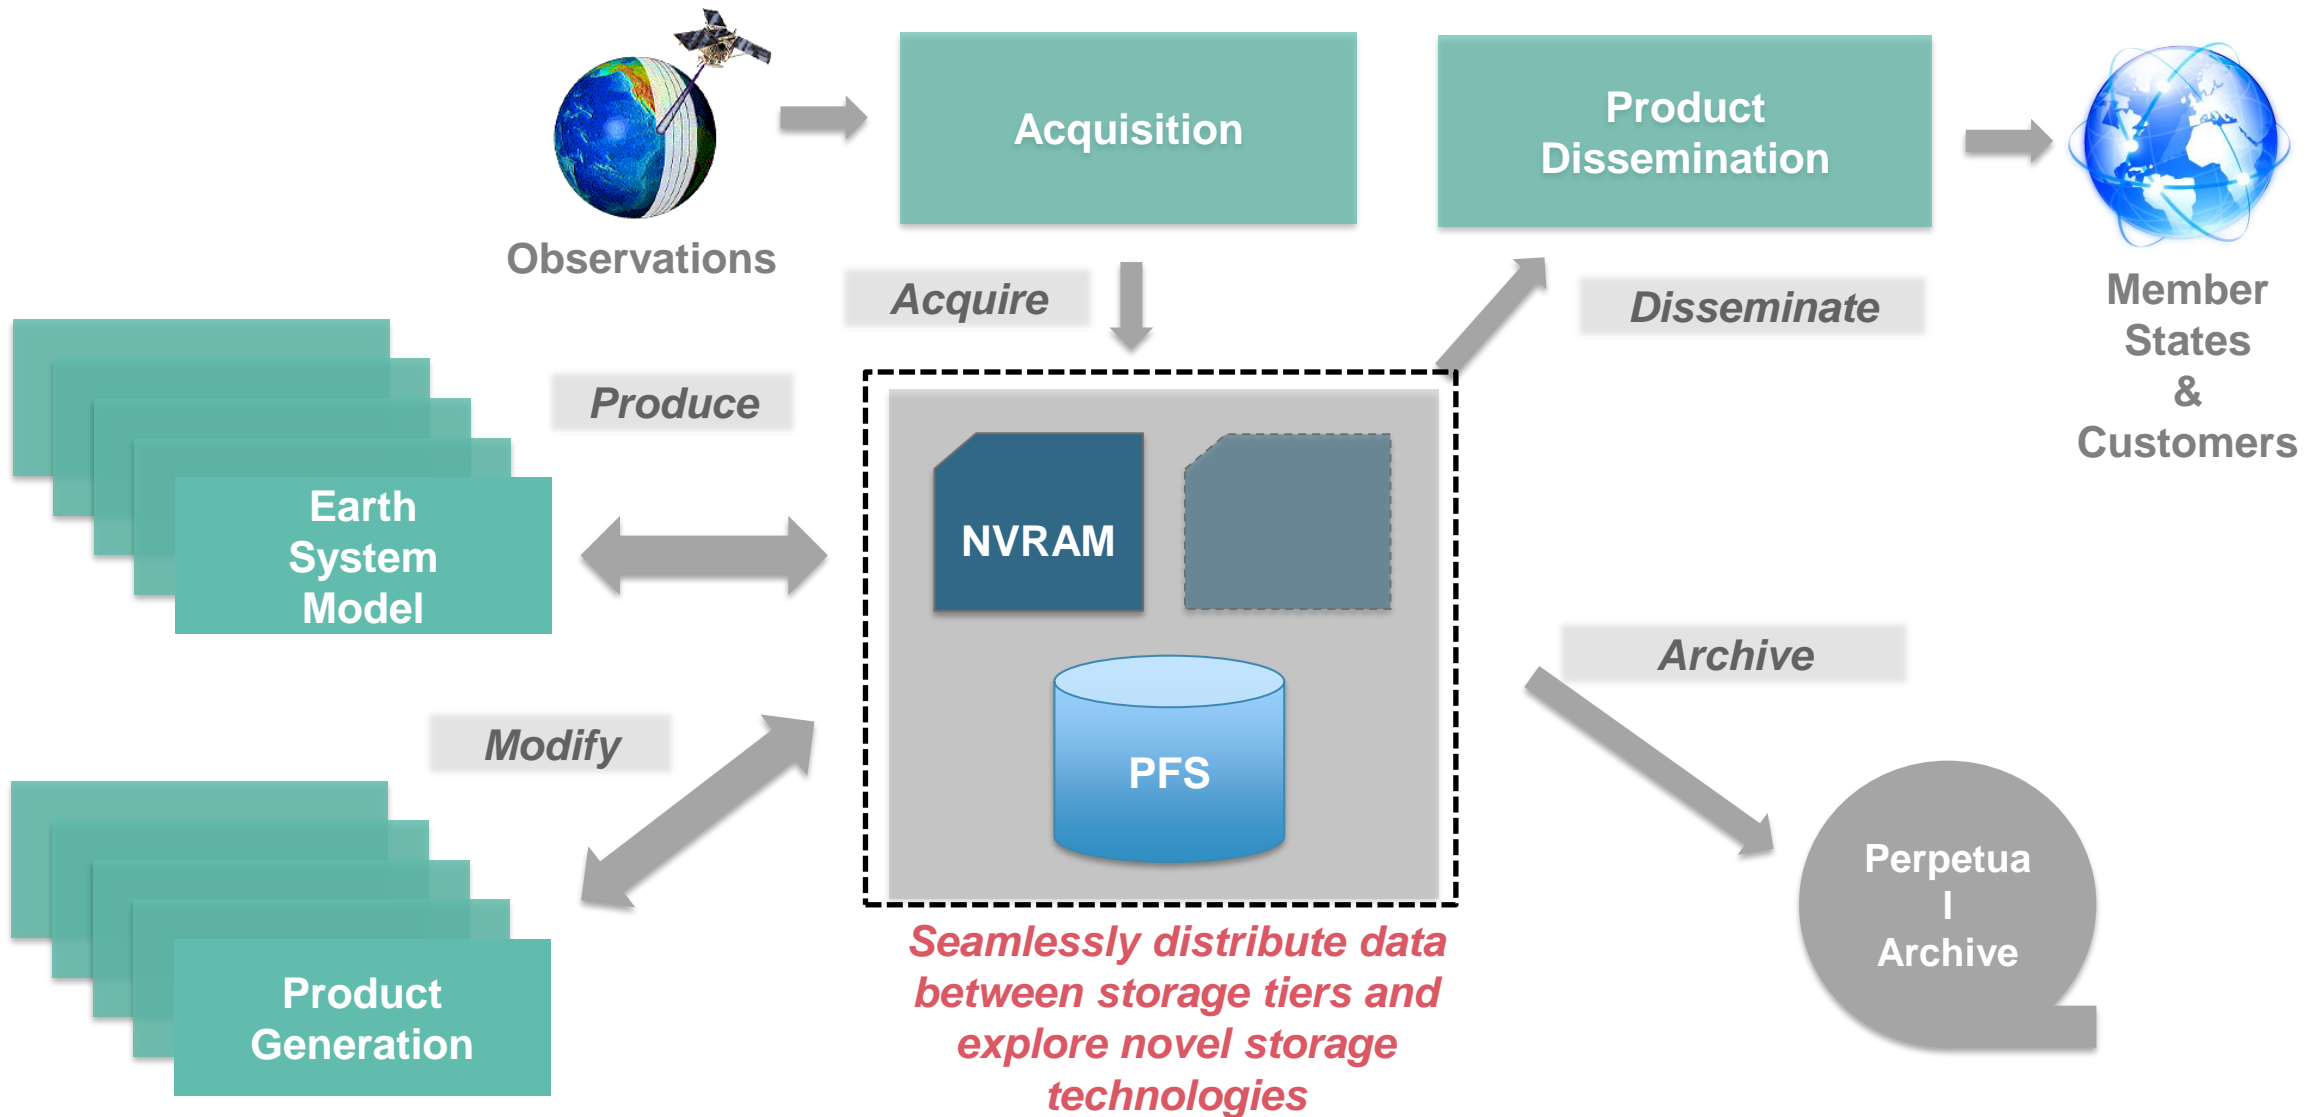
# ECMWF's Production Workflow



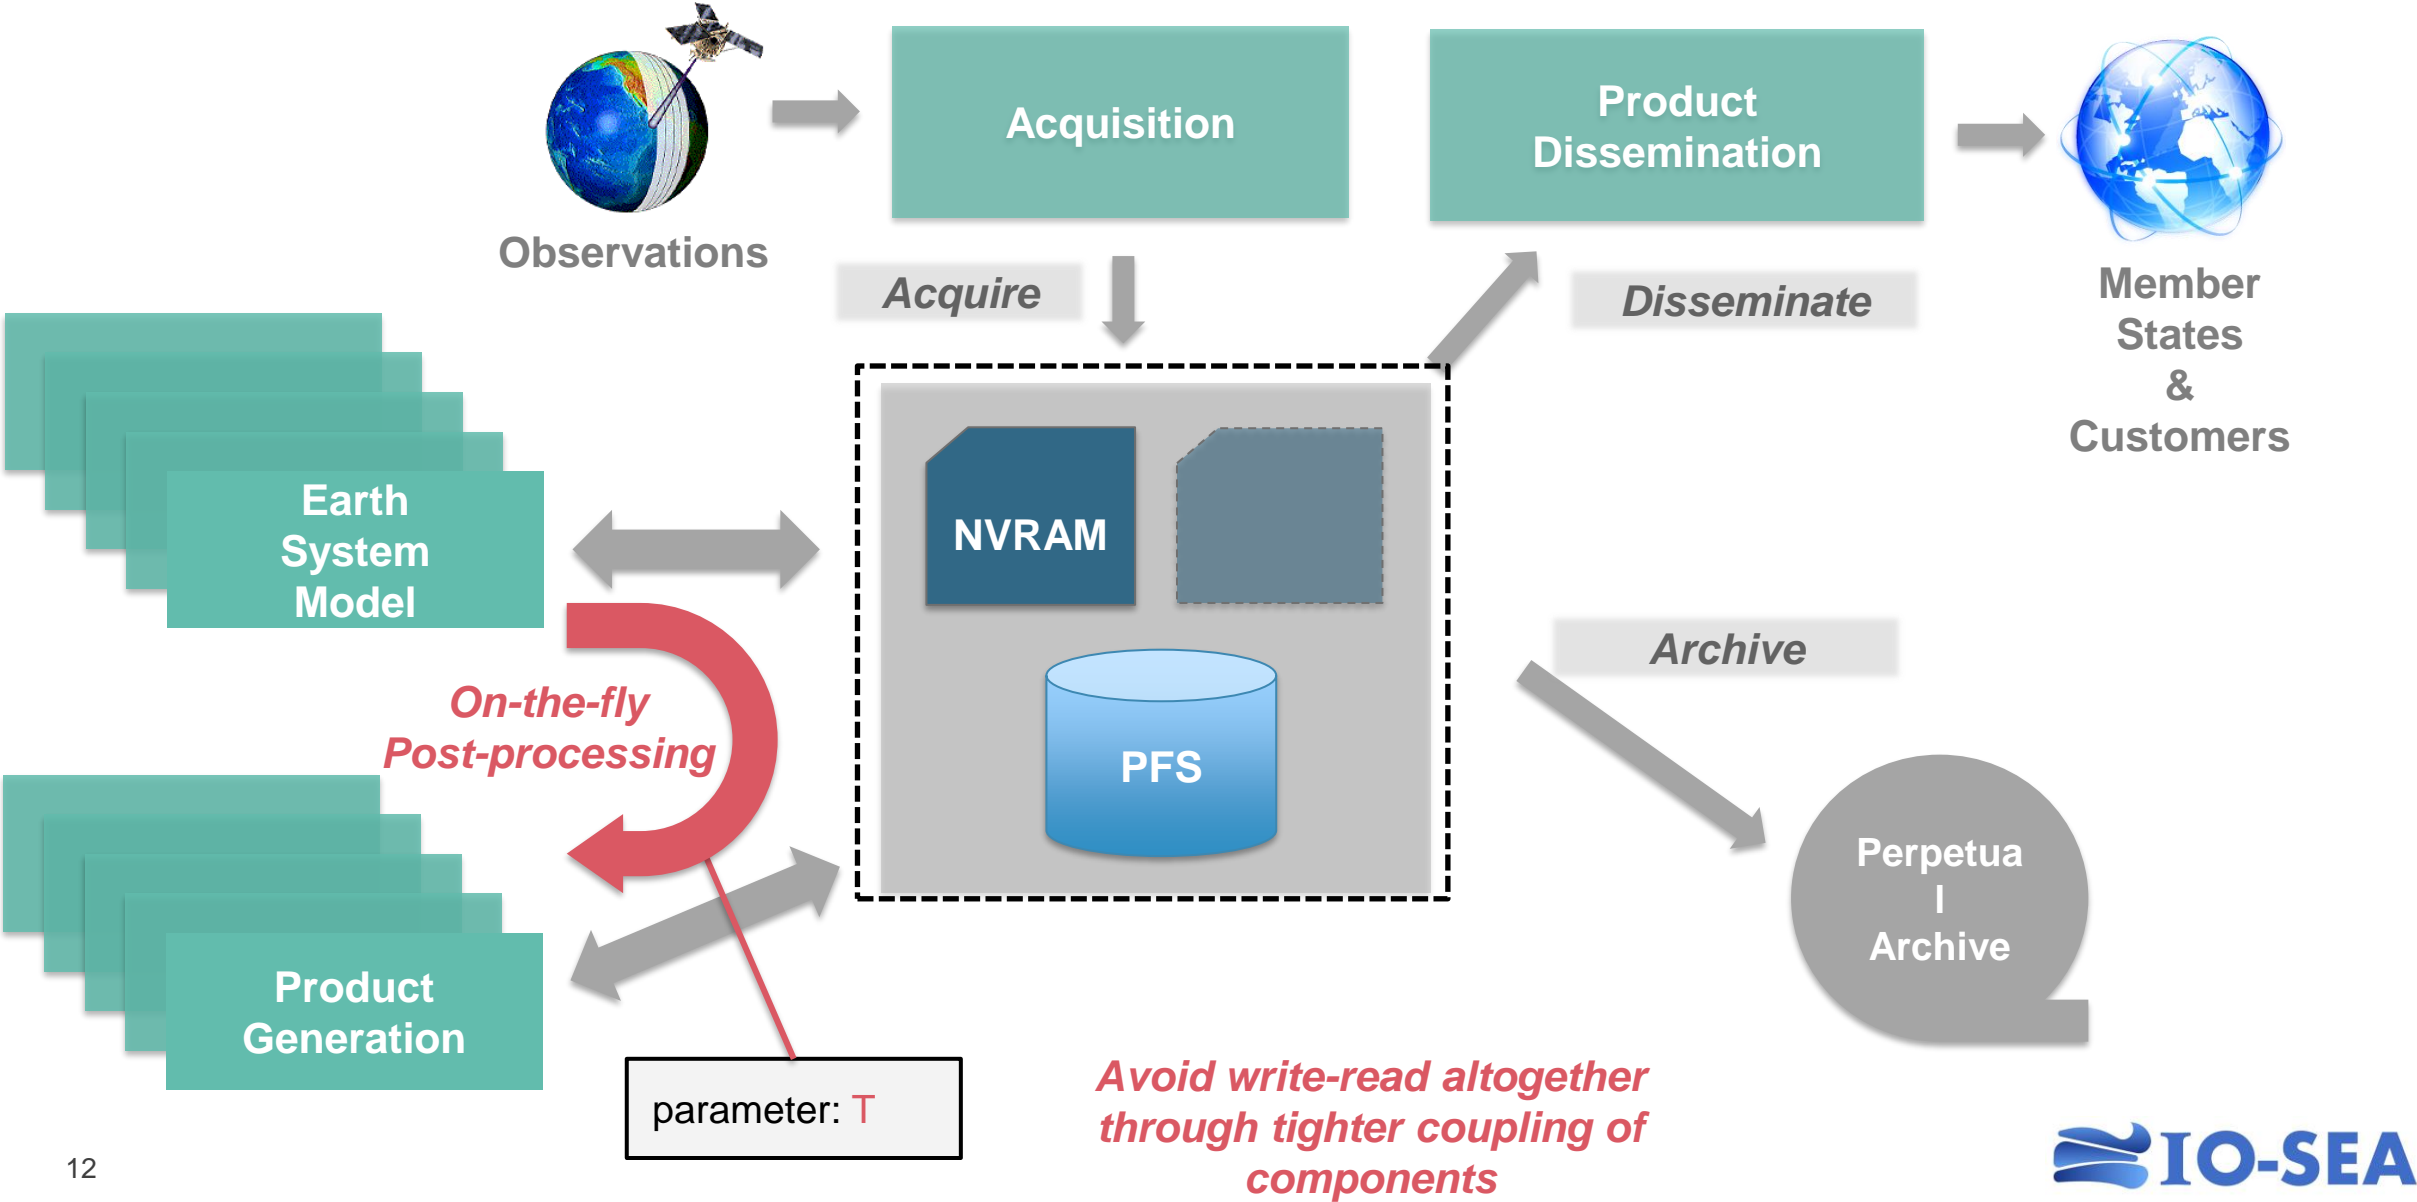
## ECMWF's Production Workflow



## Semantic Data Access > Flexible Data Storage



# Semantic Data Access > Flexible Data Routing



## ECMWF's IO Software Stack

- Data storage (FDB)
- Data aggregation and post-processing (MultIO)
- Serving and extraction of data (Polytope)
- Data notification (Aviso)
- 400 PiB tape archive (MARS)
- and others...

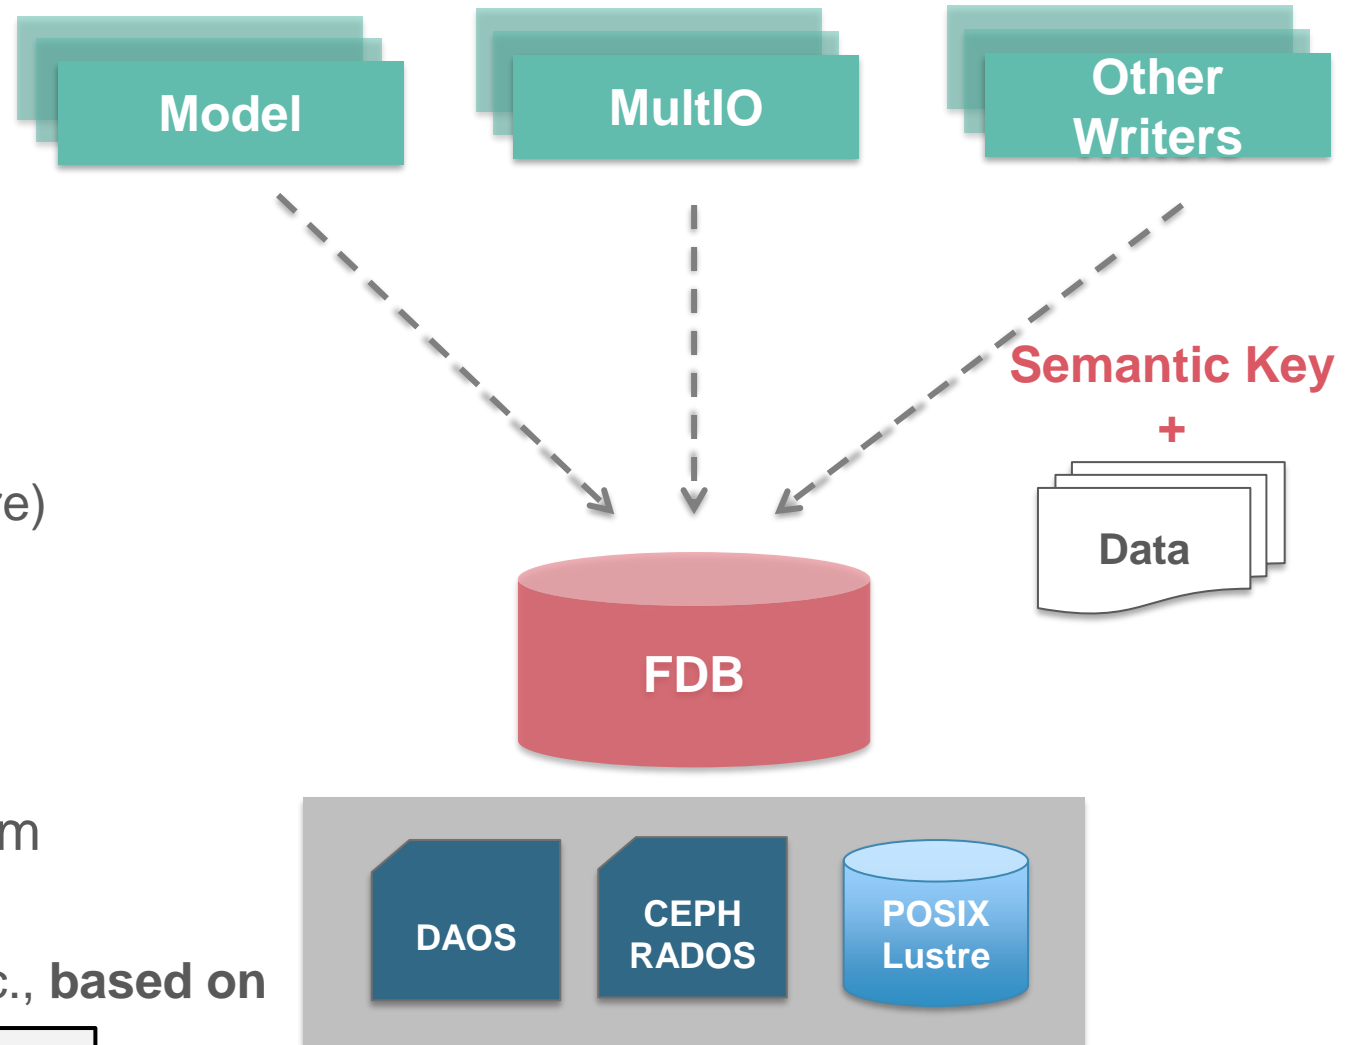
## ECMWF's IO Software Stack

- Data storage (FDB)
  - Data aggregation and post-processing (MultIO)
  - Serving and extraction of data (Polytope)
  - Data notification (Aviso)
  - 400 PiB tape archive (MARS)
  - and others...
- Component interfaces are all based on semantic data indexing using the MARS language

## FDB – Distributed Object Store

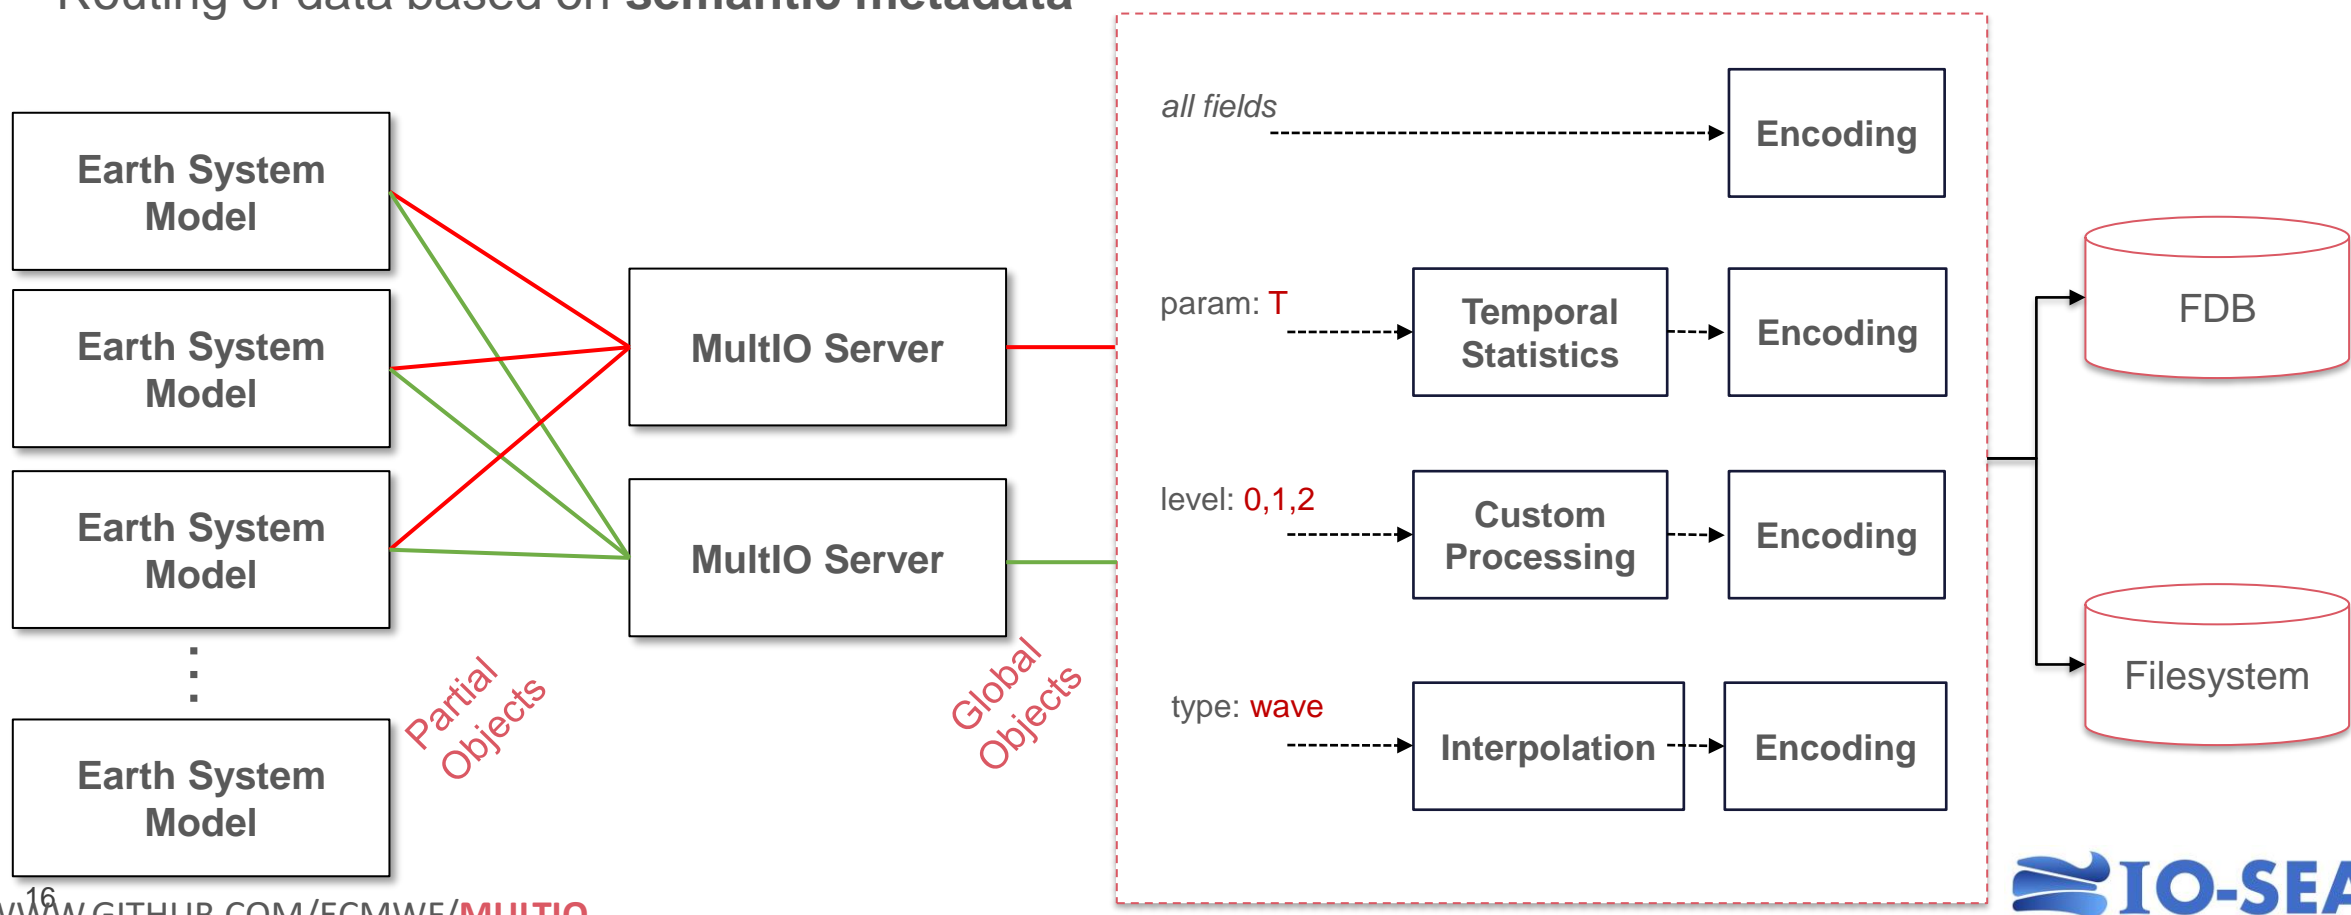
- Key-value store
  - Keys are **scientific metadata**
  - Values are byte streams
- Support for multiple back-ends
  - POSIX file-system (currently on Lustre)
  - Intel DAOS
  - CEPH/RADOS
  - Motr Cortx
- Actual storage layout abstracted from application
- Supports wild card searches, ranges, etc., **based on data semantics:**

```
date: 20210112  
time: 1200  
step: *  
level: 0 to 120
```



# MultIO – Multiplexing IO-Server

- Aggregation from distributed model
- On-the-fly post-processing
- User-defined post-processing actions
- Routing of data based on **semantic metadata**






# Polytope – Feature Extraction from Datacubes

- Hierarchical data schema defines an n-dimensional **datacube**  

`datetime x step x parameter x level x latitude x longitude`
- **Polytope** provides a way to extract data from these hypercubes
  - Flight paths, trajectories, regions, time-series, spatiotemporal polygons
- Polytope shapes defined on top of the **semantic data language**
- Exposed externally through a web API



## AVISO – Notification Service

- Notifies of data availability
- For automating systems
- Query uses **semantic keys**

**When this** → **Do that**

*data available*  
*product computed*  
*ensemble complete*


*my\_script.py*  
*FDB retrieve*  
*HTTP POST*

```
date: *  
time: *  
step: *  
parameter: T
```

```
download_data()  
produce_plot()
```

## MARS Archive

- MARS: Largest Meteorological Archive in the world
  - 400 PiB of meteorological data (tape)
  - 150 TiB added per day
- 5000 daily active users
- 2 million requests per day
- Perpetual data storage since 1985
- Every byte addressed using MARS language



# MARS Archive

- MARS: Largest Meteorological Archive in the world
  - 400 PiB of meteorological data (tape)
  - 150 TiB added per day
- 5000 daily active users
- 2 million requests per day
- Perpetual data storage since 1985
- Every byte addressed using MARS language



```
query = {  
  "class": "operational",  
  "stream": "ensemble",  
  "type": "perturbed",  
  "number": "1/to/10",  
  "date": "2023-01-10/to/2023-01-31",  
  "time": "00/12",  
  "step": "0/to/144/by/12",  
  "leveltype": "pressure",  
  "levellist": "ALL",  
  "param": "2t/p"  
}
```

# operational data  
# ... from the ensemble forecasts  
# ... specifically perturbed ensemble members  
# ... and specifically the first 10 members  
# for each forecast produced on these dates  
# ... at 00Z and 12Z UTC  
# the forecast at timestep 12, 24, etc.  
# pressure levels  
# ... the entire range of them  
# temperature and pressure parameters



# IO-SEA Hackathon

Hands-On Training with Data Access & Storage Interface (DASI)



## 13:30 Introduction:

- ECMWF and the importance of semantic data management
- IOSEA and DASI

## 14:15 Coffee Break

## 14:30 Hands on: archiving data with DASI

## 15:15 Coffee Break

## 15:30 Hands on: retrieving data with DASI

## 16:00 Q&A

**James Hawkes**  
**Metin Cakircali**  
**Jenny Wong**



**EuroHPC**  
Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955811. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, the Czech Republic, Germany, Ireland, Sweden, and the United Kingdom.