



# Easier Access to ESM Data: Implementation at Jülich Supercomputing Centre (JSC)

Carsten Hinz<sup>1</sup>, Sander Apweiler<sup>1</sup>, Simon Grasse<sup>1</sup>, Björn Hagemeier<sup>1</sup>, Rajveer Saini<sup>1</sup>, and Martin Schultz<sup>1,2</sup>


<sup>1</sup>JSC, Forschungszentrum Juelich GmbH, Juelich, Germany, <sup>2</sup>University of Cologne, Cologne, Germany





## Scientific Motivation

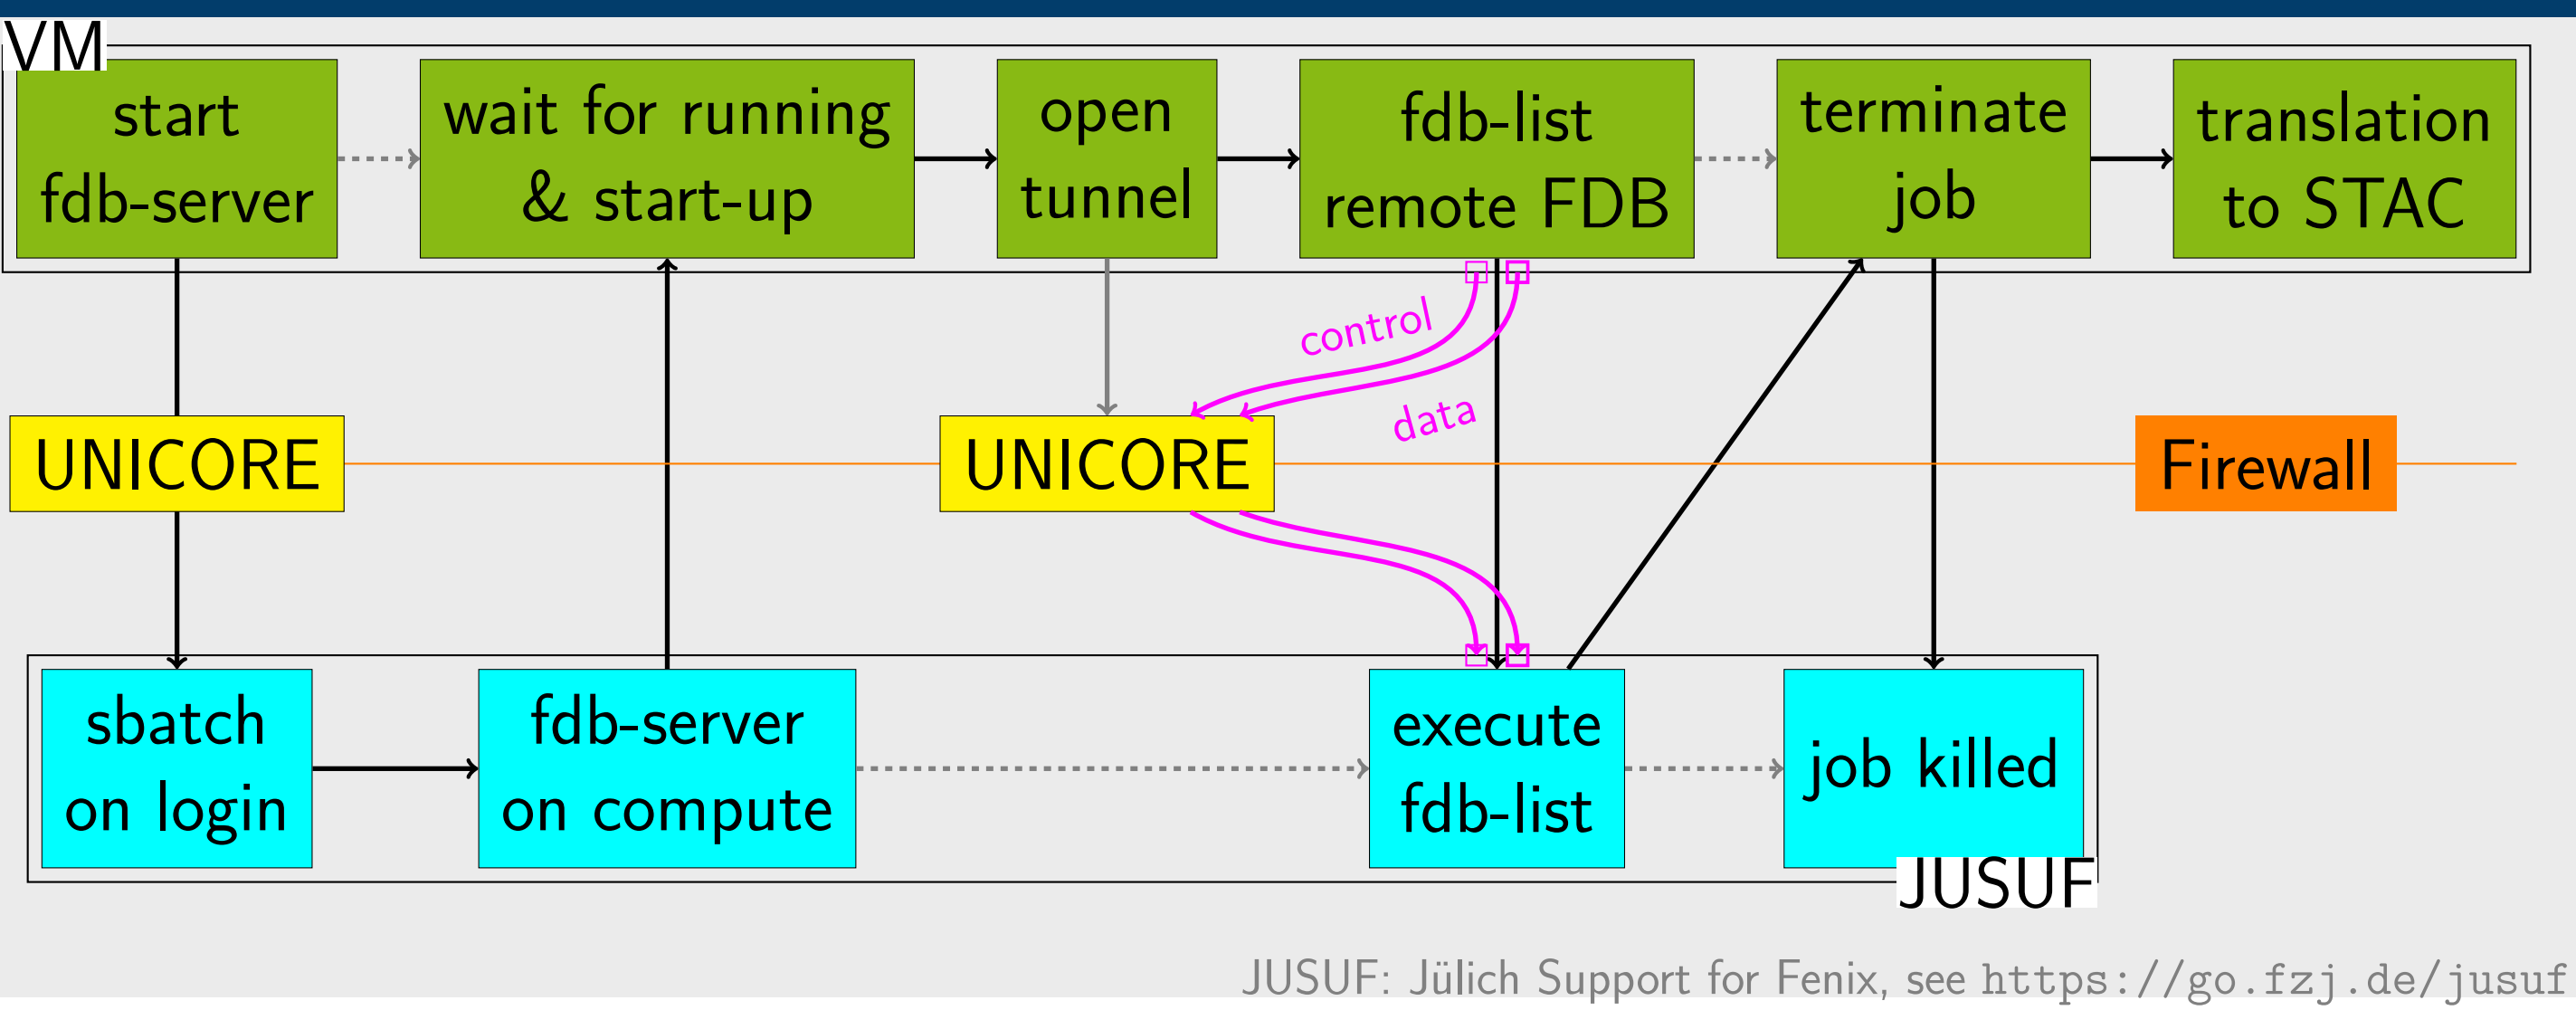
The Challenge	Easier Module within WarmWorld project	Joined Approach at DKRZ & JSC
<ul style="list-style-type: none"><li>km-scale ESM produce <math>\mathcal{O}(\text{GB})</math> per variable &amp; level<ul style="list-style-type: none"><li>Driven by advances in computing performance of HPC systems</li></ul></li><li>Challenge on storing and processing data<ul style="list-style-type: none"><li>petabyte-scale capacity file storage @ HPC centers</li><li>Prohibitive bottleneck especially for local user systems</li></ul></li></ul> <p>⇒ Demand for FAIR data, in particular findable data</p> <p>ESM=Earth System Models HPC=high performance computing FAIR=Findable, Accessible, Interoperable, and Re-usable</p>	<ul style="list-style-type: none"><li>Provide harmonized access to ICON modelling system alongside IFS-based solutions and observational<ul style="list-style-type: none"><li>Simplification of data access</li><li>Preparation to manage km-scale/exascale data streams</li></ul></li><li>Prototype exascale analysis workflows<ul style="list-style-type: none"><li>Combining simulations and observations</li></ul></li><li>Performance in WarmWorld <i>Better &amp; Faster</i> ICON=Icosahedral Nonhydrostatic[1]: a numeric ESM model by DWD, MPI-M, DKRZ, KIT, C2SM</li></ul>	<ul style="list-style-type: none"><li>Federated catalog for data centers</li><li>Common interfaces to access data &amp; analysis e.g. direct/local access, download, zarr over HTTP</li><li>Easier use-case drives developments:<ul style="list-style-type: none"><li>hosting of IFS data @ JSC and ICON data @ DKRZ for analysis by University of Cologne</li></ul></li></ul> <p>IFS=Integrated Forecast System of ECMWF[2] DKRZ=German Climate Computing Center ECMWF=European Centre for Medium-Range Weather Forecasts</p>

## Ingredients for developments at JSC

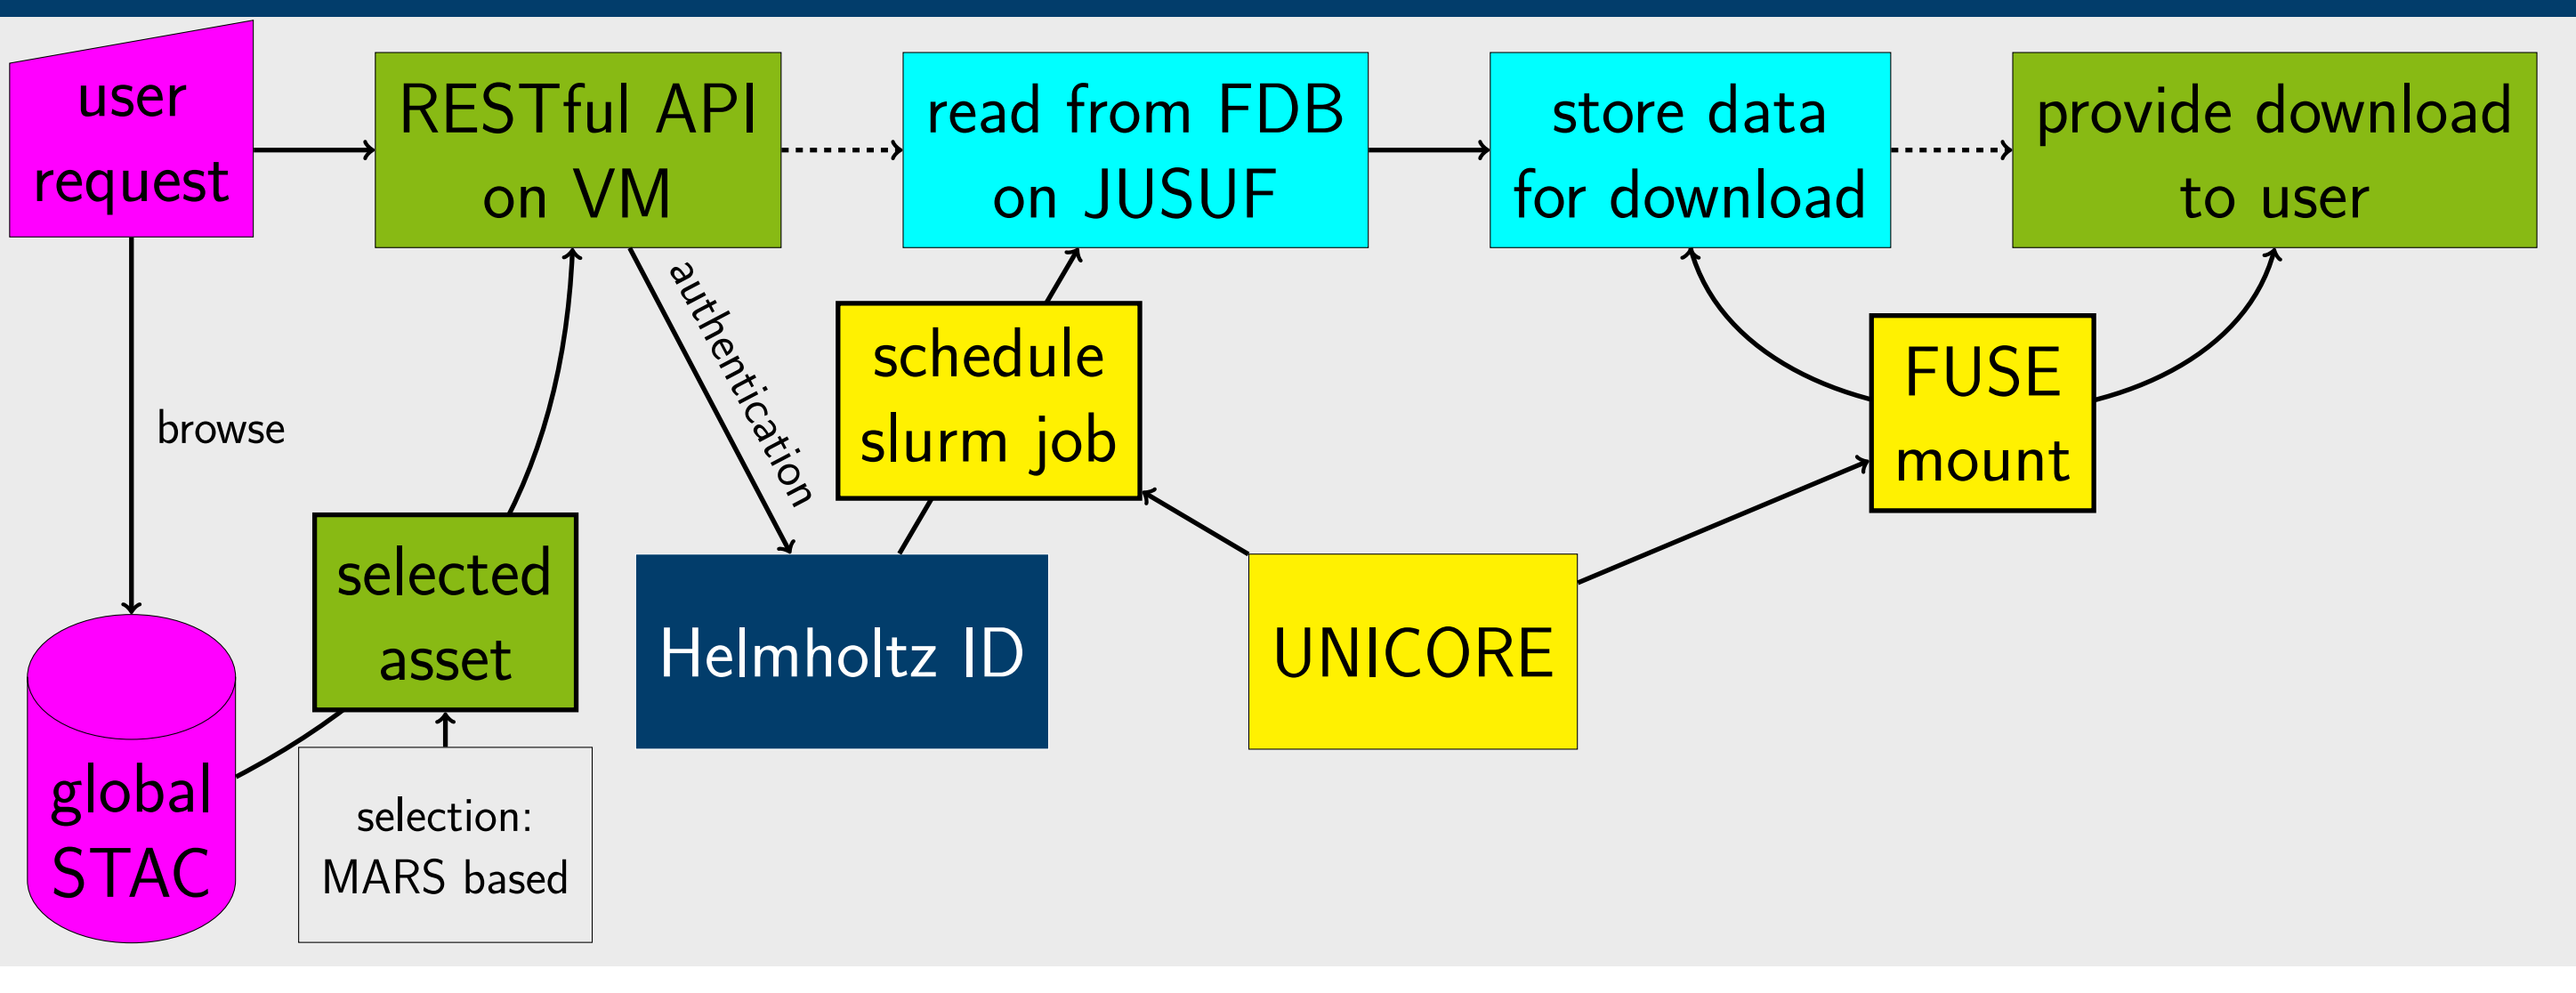
Data Backend: Fields DataBase (FDB)	SpatioTemporal Asset Catalog (STAC)	Test Dataset
Object store for geospatial data[3, 4] <ul style="list-style-type: none"><li>Operational storage for recent results at ECMWF</li><li>Access to data via metadata (MARS language)</li></ul> e.g. Storage of data in POSIX structure <ul style="list-style-type: none"><li>Directory → file → grib file (multiple fields)</li><li>Structure defined by metadata (configurable)</li></ul>	Metadata catalog for geospatial data[5] <ul style="list-style-type: none"><li>Based on JSON: Machine and human readable</li><li>JSC: Hosting API specification with Postgres backend[6]</li><li>Federated catalog hosted at DKRZ</li><li>More details: Kameswar Rao Modali et al. (EGU25-10288 at Hall X4   X4.107)</li></ul> 	IFS dataset (306 600 fields) <ul style="list-style-type: none"><li>365 days with two runs/day and 12 forecast steps</li><li>5 model levels</li><li>7 atmospheric and 1 surface variable</li><li>0.25° × 0.25° resolution (0.25° ≈ 27 km)</li><li>≈2 MB per field</li></ul>

UNICORE - UNiform Interface to COmputing REsources	Helmholtz ID as Authorization and Authentication Infrastructure
Access to HPC resources via RESTful API through different interfaces[7] <ul style="list-style-type: none"><li>Start and control of jobs on HPC systems (e.g. SLURM)</li><li>Port forwarding from external systems to compute nodes</li><li>Access to storage cluster, e.g. FUSE mounts</li><li>Allows data driven execution of scripts</li></ul> 	OpenID Connect (OIDC) provider for Helmholtz centers and other scientific institutions[8] <ul style="list-style-type: none"><li>Federated AAI also including social Identity Providers (ORCiD, github, google)</li><li>Login through infrastructure of identity provider/home institution</li><li>Authorization via different scopes possible, e.g. handling of "virtual organisations"</li></ul> 

## Recipe for Filling a STAC Catalog with Metadata from an FDB

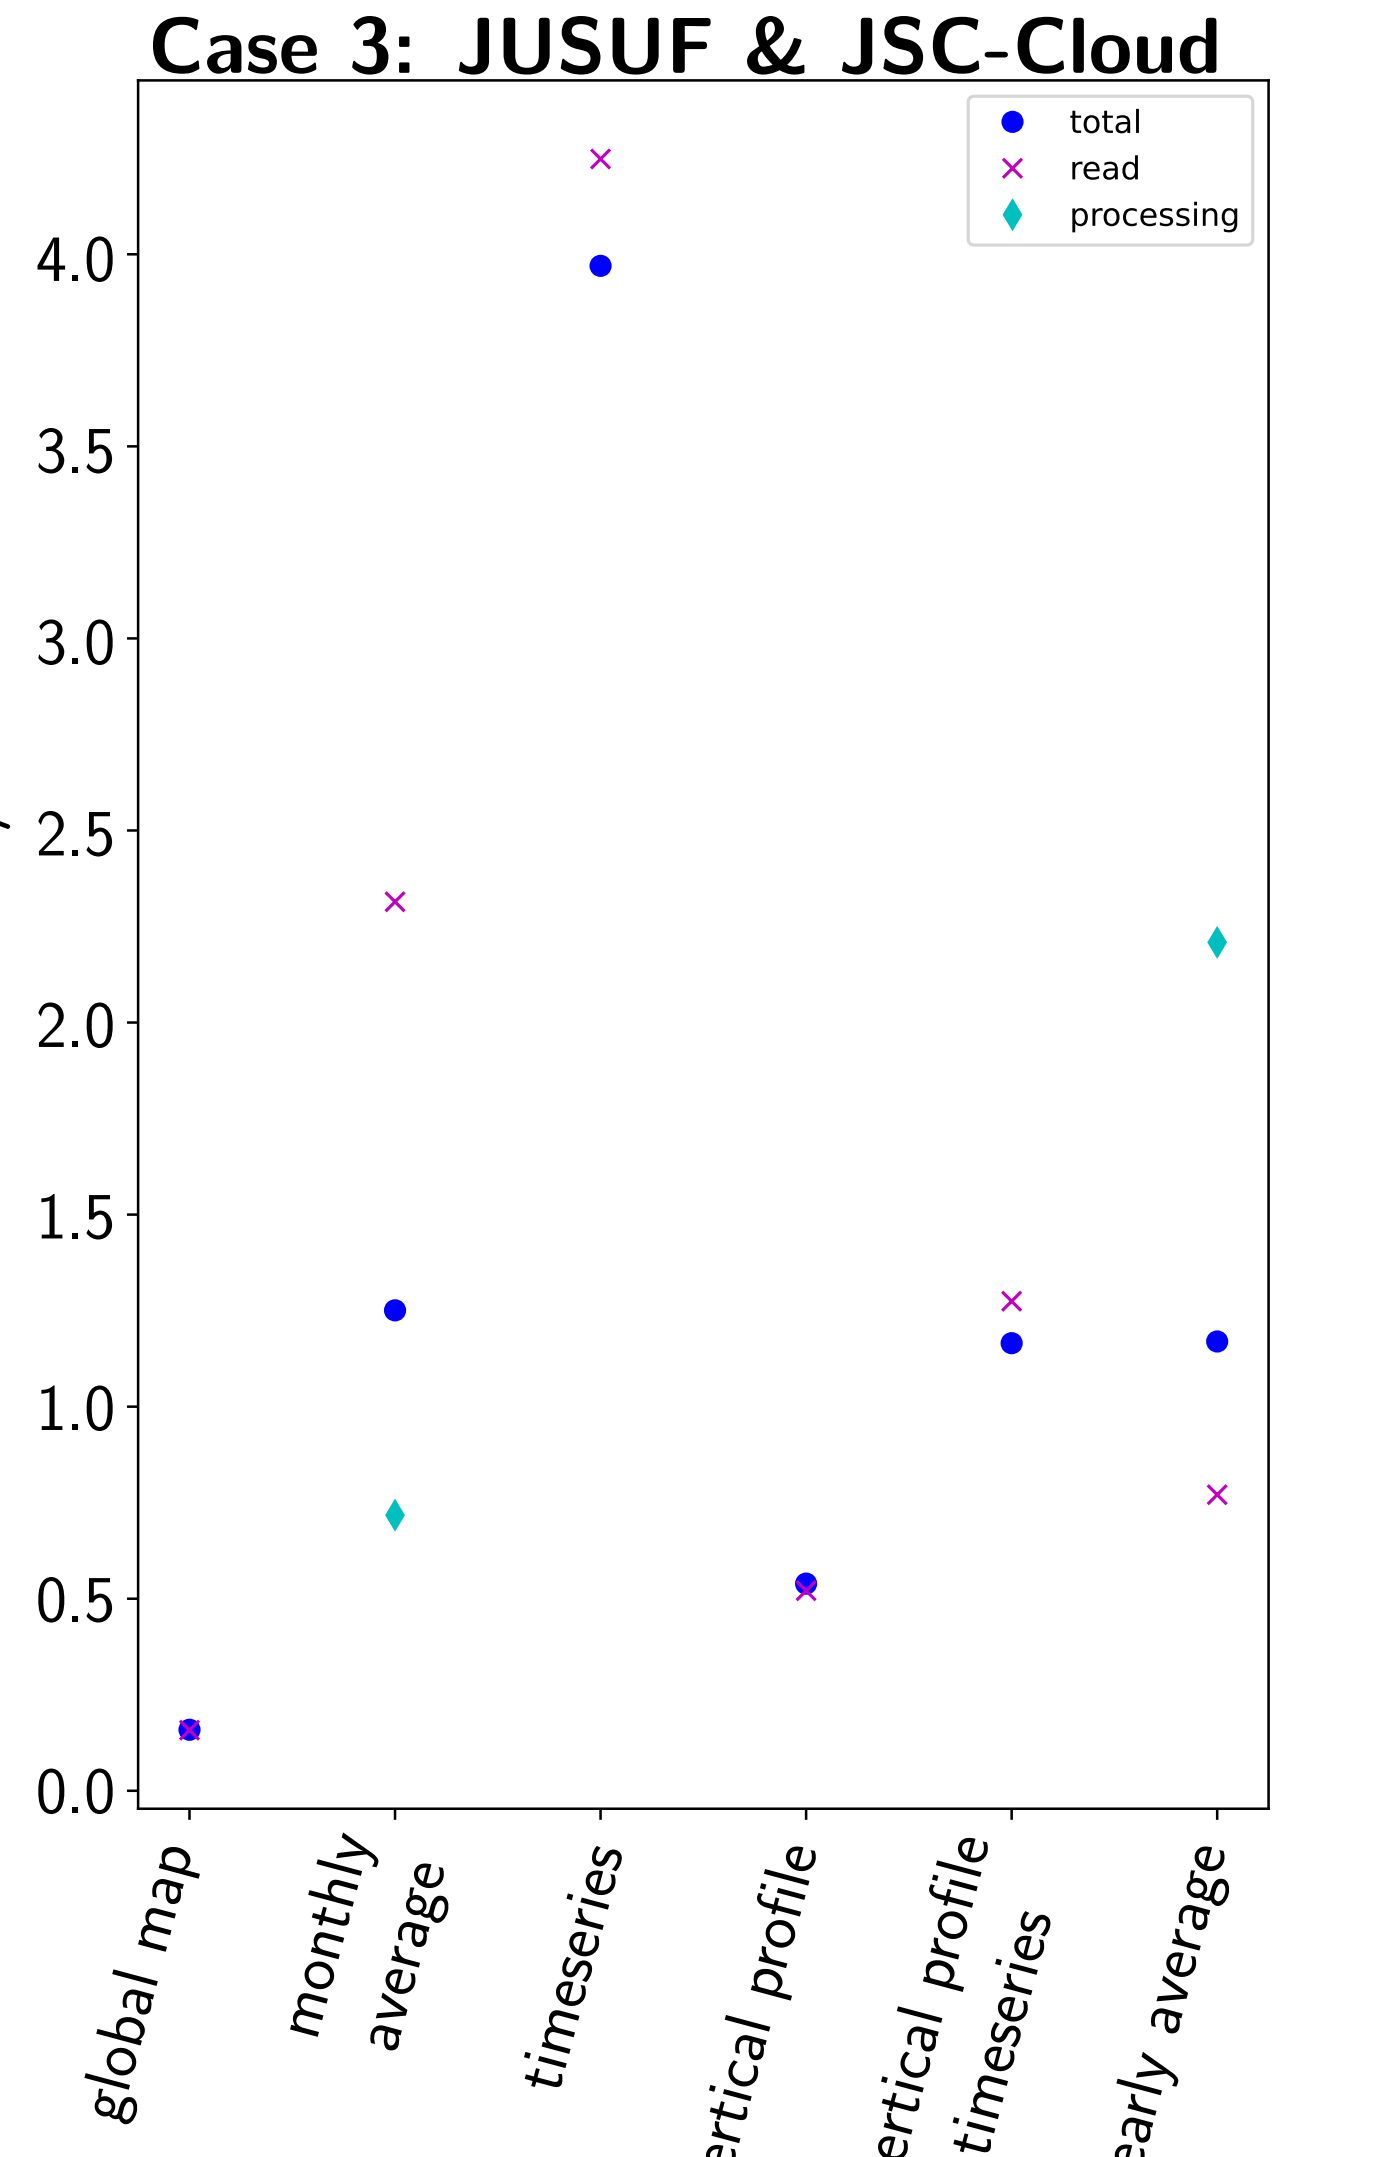
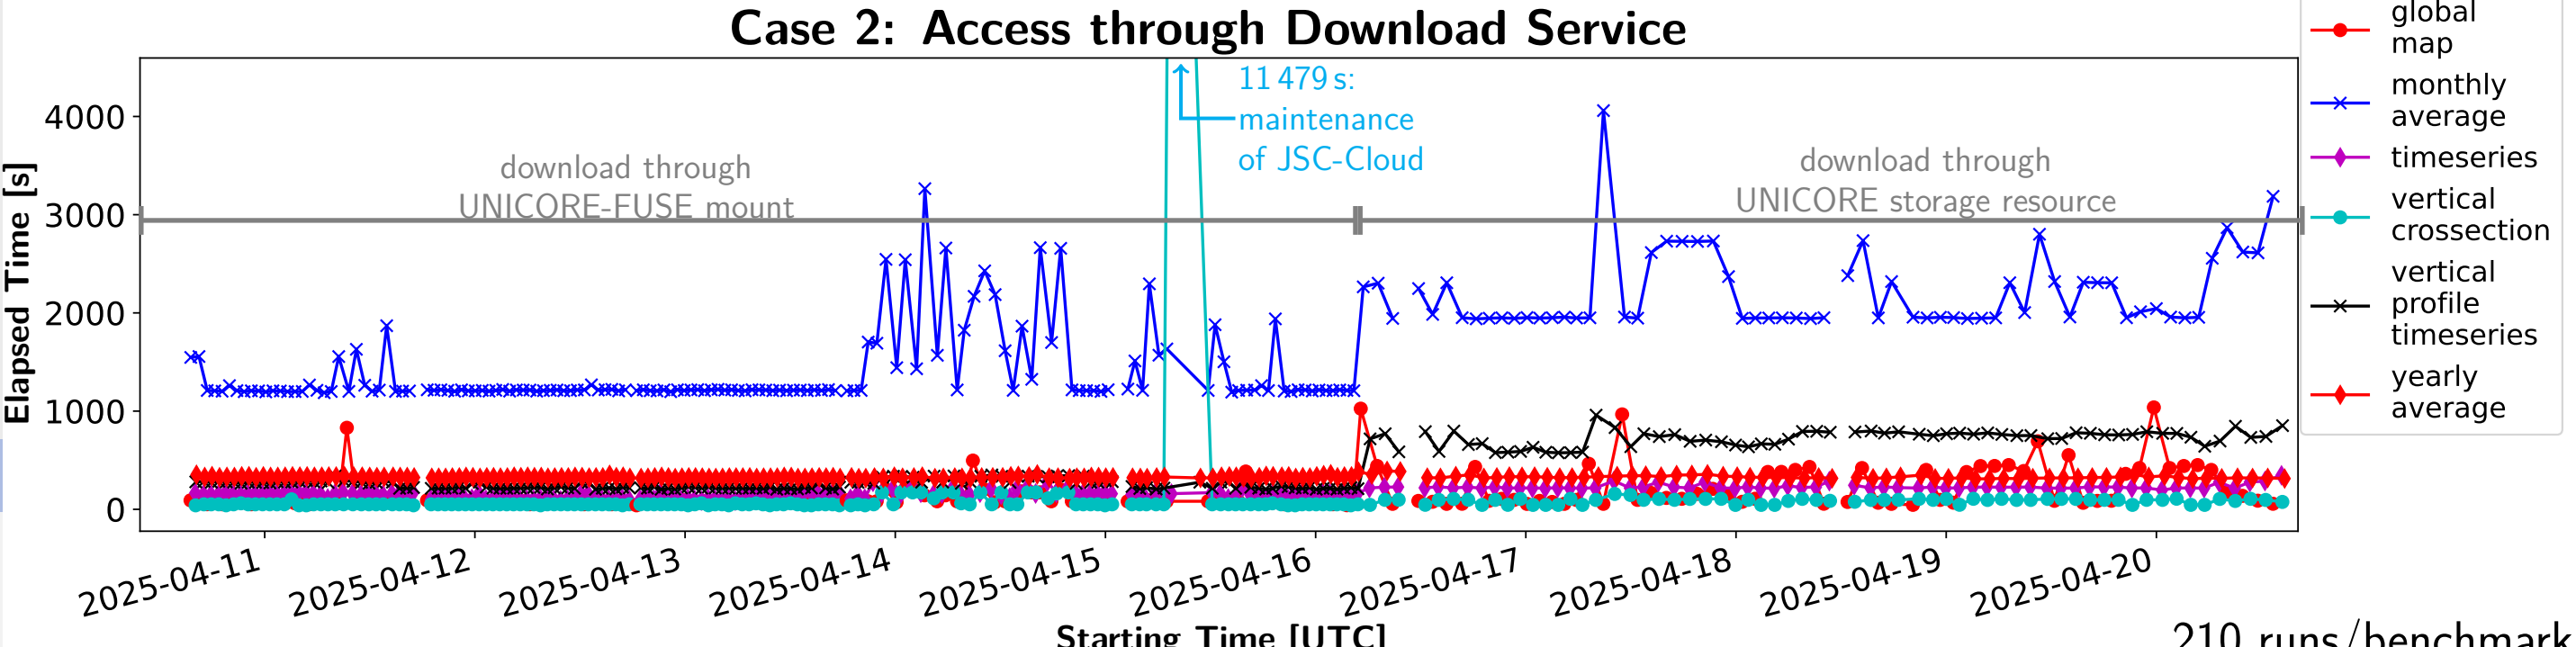
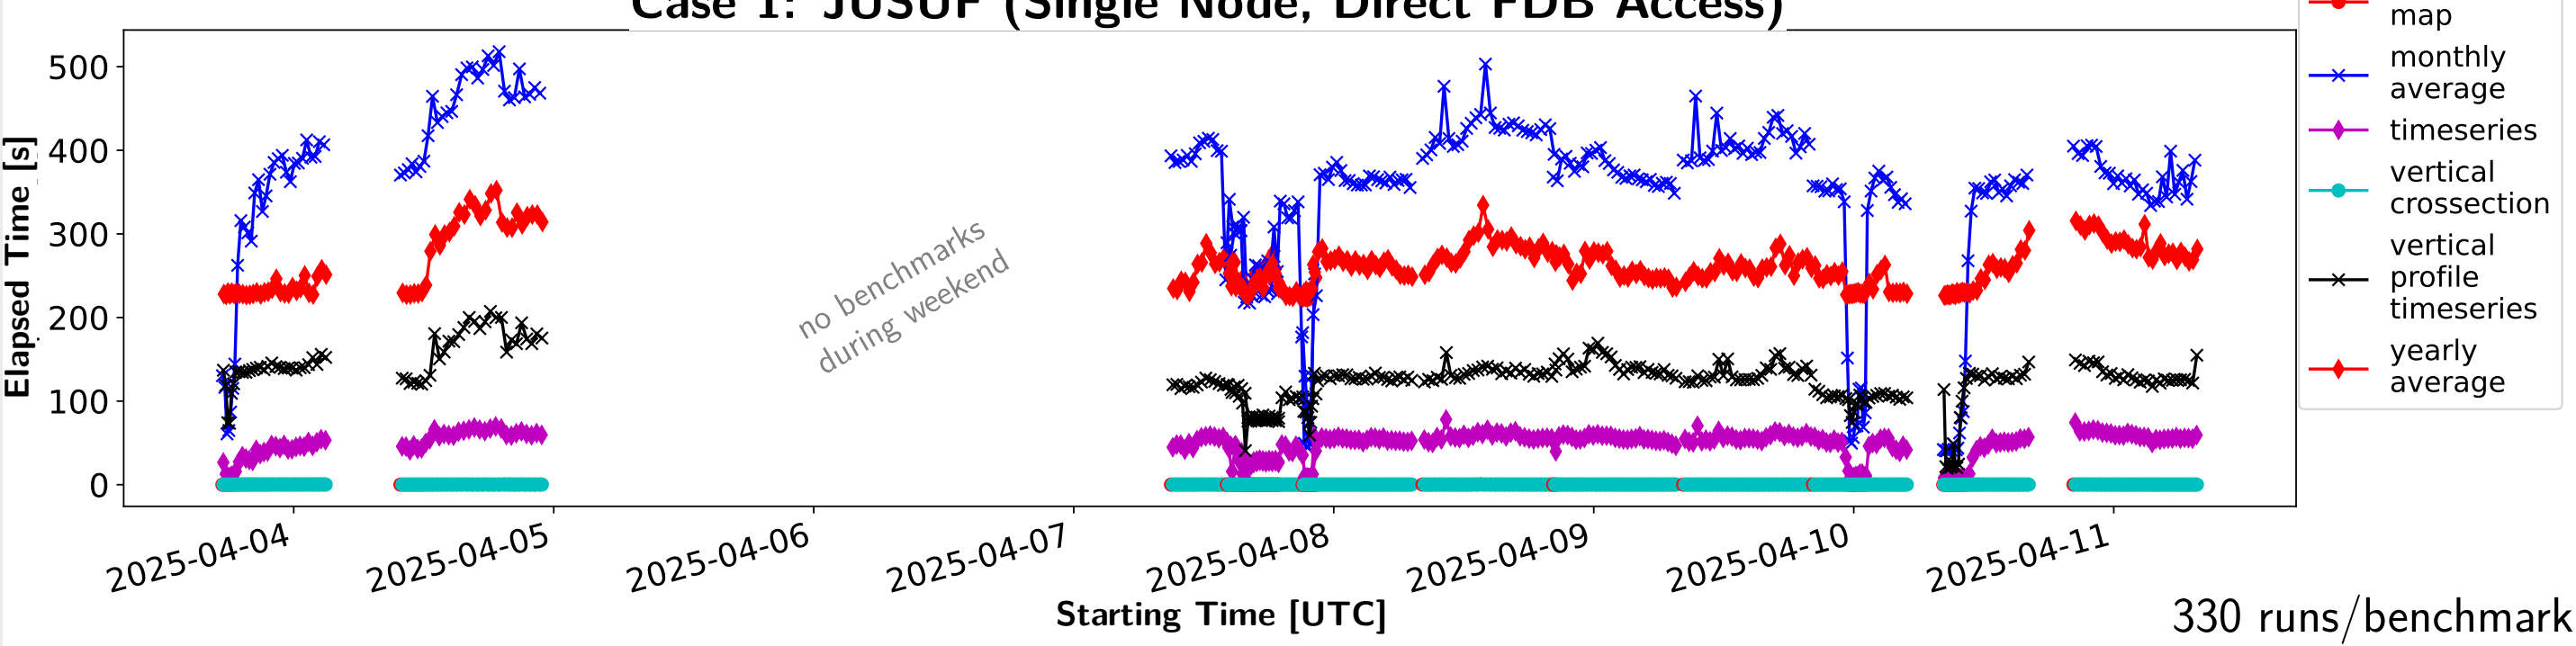


## Recipe for Download Service



## Outcome: Benchmarks for Accessing Data in FDB

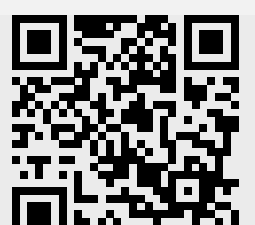
- Case 1: Local operation on JUSUF, direct access to FDB
  - JUSUF as HPC system has access to shared storage cluster JUST hosting the FDB
  - FDB @ \$SCRATCH
  - Recent simulations will be written \$SCRATCH
- Case 2: Remote user with access to FDB through JSC-Cloud
  - No direct access to FDB
  - Cloud has no access to \$SCRATCH
  - Cloud hosts services to download data
- Case 3: Local operation on VM
  - FDB on dedicated volume
- Cloud as host for data management services (e.g. fill catalog)



Benchmark	global map	monthly average	timeseries	vertical profile	vertical profile timeseries	yearly average
Fields	1	8760	730	5	3650	8760
Median FDB-JUSUF	0.154 s	365.144 s	52.891 s	0.205 s	127.233 s	255.102 s
Median DL-Service JUSUF	82.022 s	1218.557 s	160.136 s	51.078 s	276.338 s	325.017 s
Median FDB-VM	0.973 s	157.770 s	12.450 s	0.394 s	99.829 s	330.971 s

Hardware

	VM @ JSC-Cloud	JUSUF
CPU	4 cores@2.25 GHz	2× AMD EPYC 7742 64 cores @2.25 GHz
RAM	7.75 GB	16 × 16 GB DDR4 @ 3200 MHz
Volume	15 TB	access to JUST $\mathcal{O}(\text{TB})$ to $\mathcal{O}(\text{PB})$

JUST: Jülich Storage Cluster  
\$SCRATCH fast access but limited lifetime storage tier  
JUST specification for details: 

## Conclusions

- Storage cluster and HPC provide a compatible performance to VM
  - ! Data size requires storage cluster like JUST for km-scale simulations
- Access through service is slower (job allocation, copying on HPC system & download)
- Access for project partners through Helmholtz ID

## Outlook

- Improved data access services
  - Hosting of a service to access FDB data as zarr via HTTP (development by ECMWF)
  - Hosting of a *polytope*[9] service
  - Optimization of data extraction for download & job allocation
- Extended authorization using a "Virtual organisation" concept of the Helmholtz ID
- Portal for ESM data services provided by JSC within WarmWorld

## Acknowledgements

We thank our colleagues and project partners at ECMWF, especially Emanuele Danovaro, Tobias Kremer and Kai Kratz, for recent discussion and support on optimizing our FDB setup. We also thank our colleagues at DKRZ, especially Kameswar Rao Modali and Karsten Peters-von-Gehen, for the discussions and joined work on the federated data catalog. We would also like to thank our colleagues Sabine Schröder and Max Lensing for the discussions on the implementation of the services. The authors would like to thank the German Federal Ministry of Education and Research for the funding of this work in the scope of WarmWorld: Easier (project identification number: 01LK2204D).

Contact: c.hinz@fz-juelich.de - Website: www.fz-juelich.de/en/ias/jsc/

GEFÖRDERT VOM



Bundesministerium  
für Bildung  
und Forschung



Sharing is  
encouraged