Creation of a Public Dataspace for Earth System Data at the Jülich Supercomputing Centre

Carsten Hinz[1], Sabine Griessbach[1], Lars Hoffmann[1], Enxhi Kreshpa[1], Kameswar Rao Modali[2], Karsten Peters-von Gehlen[2], Konstantin Rushchanskii[1], Rajveer Saini[1], Olaf Stein[1], and Martin Schultz[1,3]

[1]Juelich Supercomputing Centre (JSC), Forschungszentrum Juelich GmbH, Juelich, Germany, [2]German Climate Computing Center (DKRZ), Hamburg, Germany, [3]University of Cologne, Cologne, Germany

Scientific Motivation

- Numerical km-scale climate modeling produces $\mathcal{O}$(PB) per run
- Machine learning (ML)-training requires fast read access to large volumes of data
- Only HPC offers required storage and compute capacities at scale

⇒ Increasing demand for FAIR and easily accessible data within the Earth system modeling (ESM) community

⇒ Need for a shared data infrastructure for efficient utilization

HPC Systems and Storage @ JSC and DKRZ

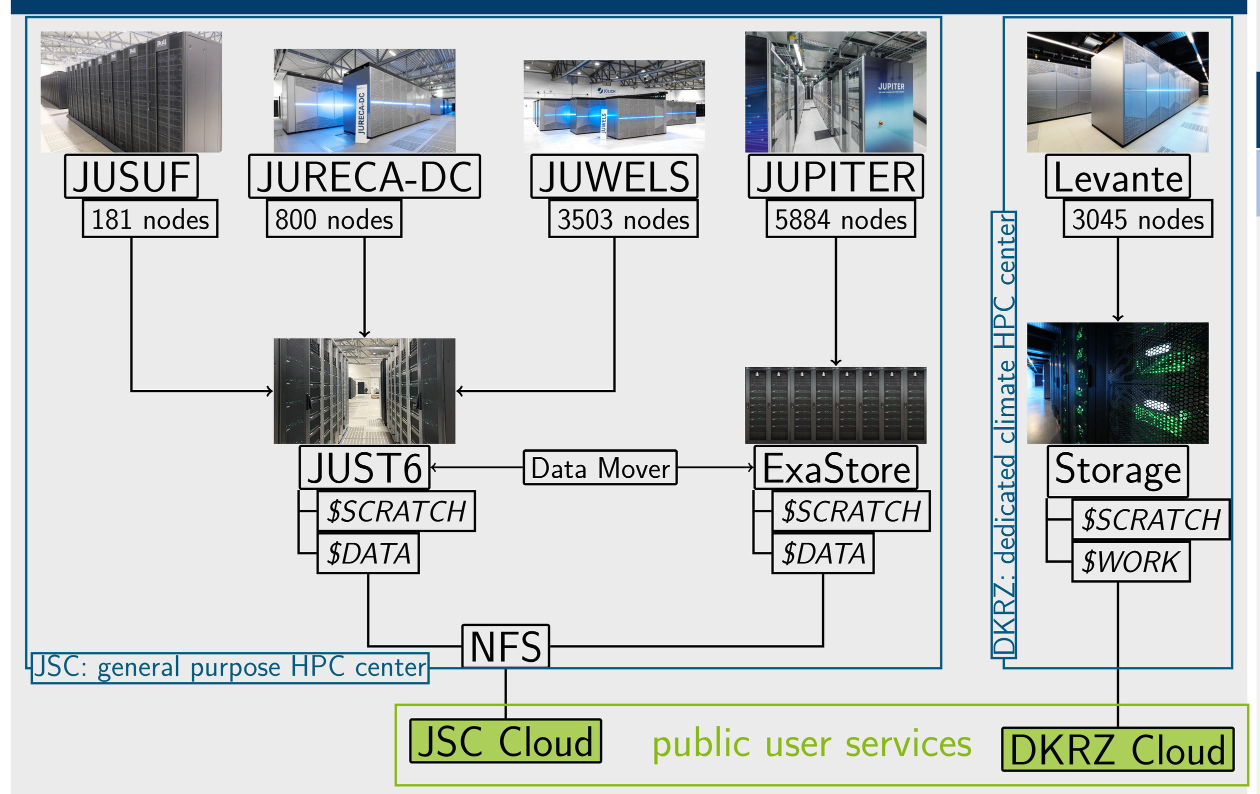

Storage Systems:

	gross capacity	capacity $DATA/$WORK	tape capacity	Center
JUST6	216.2 PB (GPFS)	49 PB	478 PB	JSC
ExaStore	> 300 PB (GPFS)	100 PB	> 700 PB	JSC
Levante	130 PB (lustre)	≈ 120 PB	300 PB	DKRZ

Numbers from JSC/DKRZ user documentation as of April 2026
Images of JSC systems originate from Forschungszentrum Juelich GmbH
Images of DKRZ systems originate from German Climate Computing Center

Background: Anemoi-Datasets

Open-source format for machine learning developed by ECMWF[1]

- Based on Zarr v2 directory store with a single array for all data
  - Gridded datasets (e.g. for model fields)
  - Tabular datasets for spatial and/or temporally unstructured data (e.g. observations)
- Array layout optimized for I/O during ML-training
  - Chunking along first axis: time
- Built from recipes to allow recreation of datasets

Background: Fields DataBase (FDB)

Object store for geospatial data[2]

- Operational storage for recent results at ECMWF
- Access to data via metadata (MARS language)
- Storage of data in POSIX structure based on metadata
- Technical access through API (e.g. C++, python) or CLI tools

Background: SpatioTemporal Asset Catalog (STAC)

Metadata catalog for geospatial data[3]

- Based on JSON: Machine and human readable
- Specification for API driven catalogs
  - Database backends like Elasticsearch or Postgres
- Customizable through extensions, e.g. Datacube

Unifying Approaches

Jülich MeteoCloud

- Shared data project for ESM community at FZJ
- Serving commonly used datasets to avoid duplicate data
- Meteorological data: e.g. ERA5, CERRA, MERRA-2
- Model Data: input data for ICON & MESSy and results for CAMS
- Satellite Data: e.g. AIRS, IASI, MIPAS
- new ML-training ready datasets for WeatherGenerator e.g. ERA5, CERRA, SEVIRI

full list: `https://datapub.fz-juelich.de/slcs/meteocloud/`

Data Access Strategy in WarmWorld-Easier Project

- Provide harmonized access to ICON modeling results alongside IFS-based solutions (ERA5, forecasts) and observations
- Preparation for km-scale/exascale data streams

⇒ Federated Data Systems Across JSC and DKRZ:
  - Unified data access through Zarr over HTTP
  - STAC for data discovery
  - Exploration of Fields Database as storage backend

project website: `https://www.warmworld.de/easier/`

Jülich MeteoCloud as a Joint Dataspace

Jülich MeteoCloud in Numbers

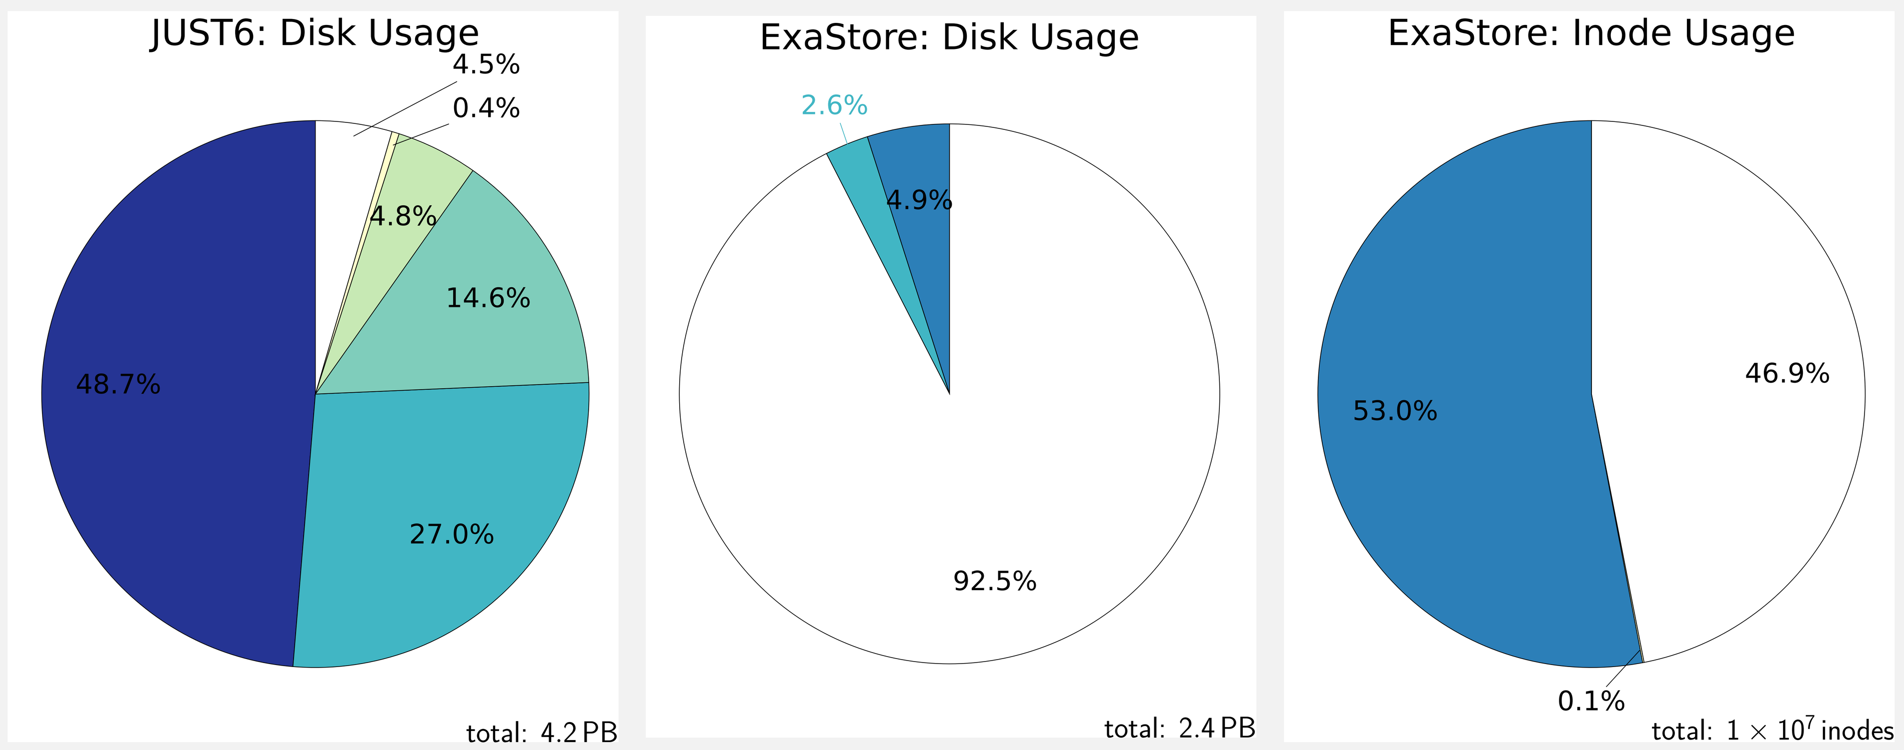

Data Access Examples: MERRA-2

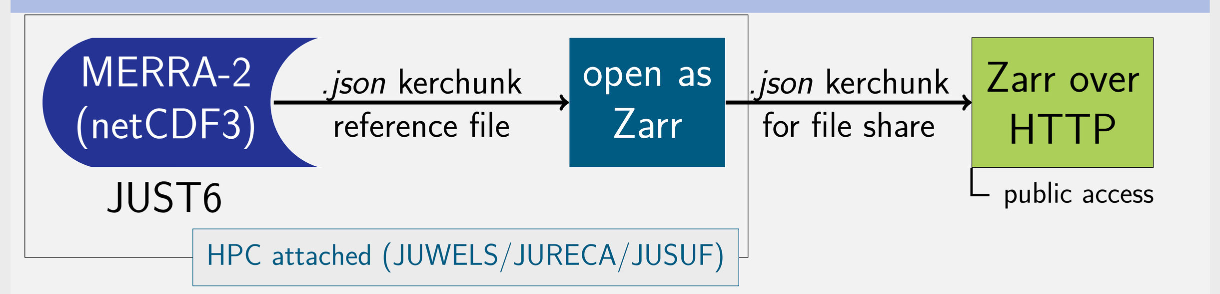

Data Access Examples: ML-training ready

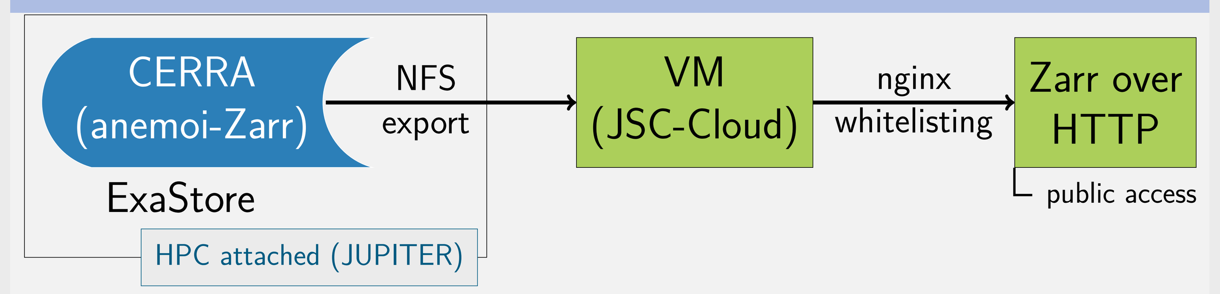

FDB as Storage Backend for GRIB & Data Access

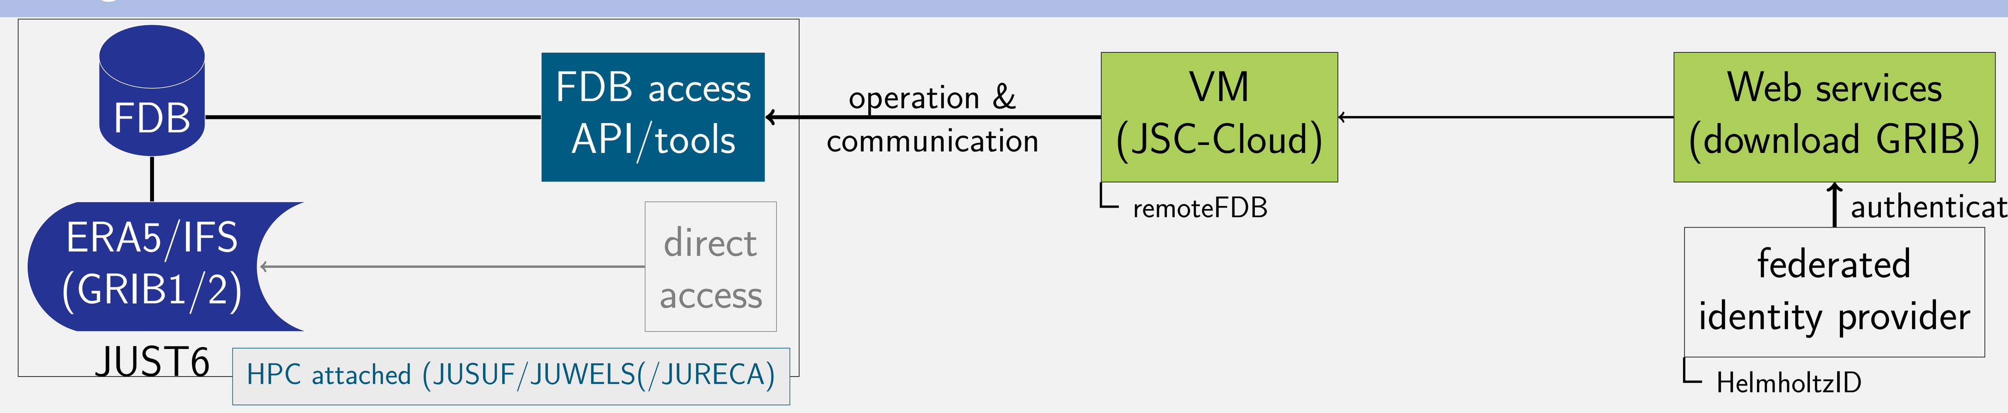

SpatioTemporal Asset Catalog to Enhance Findability with Unified Metadata

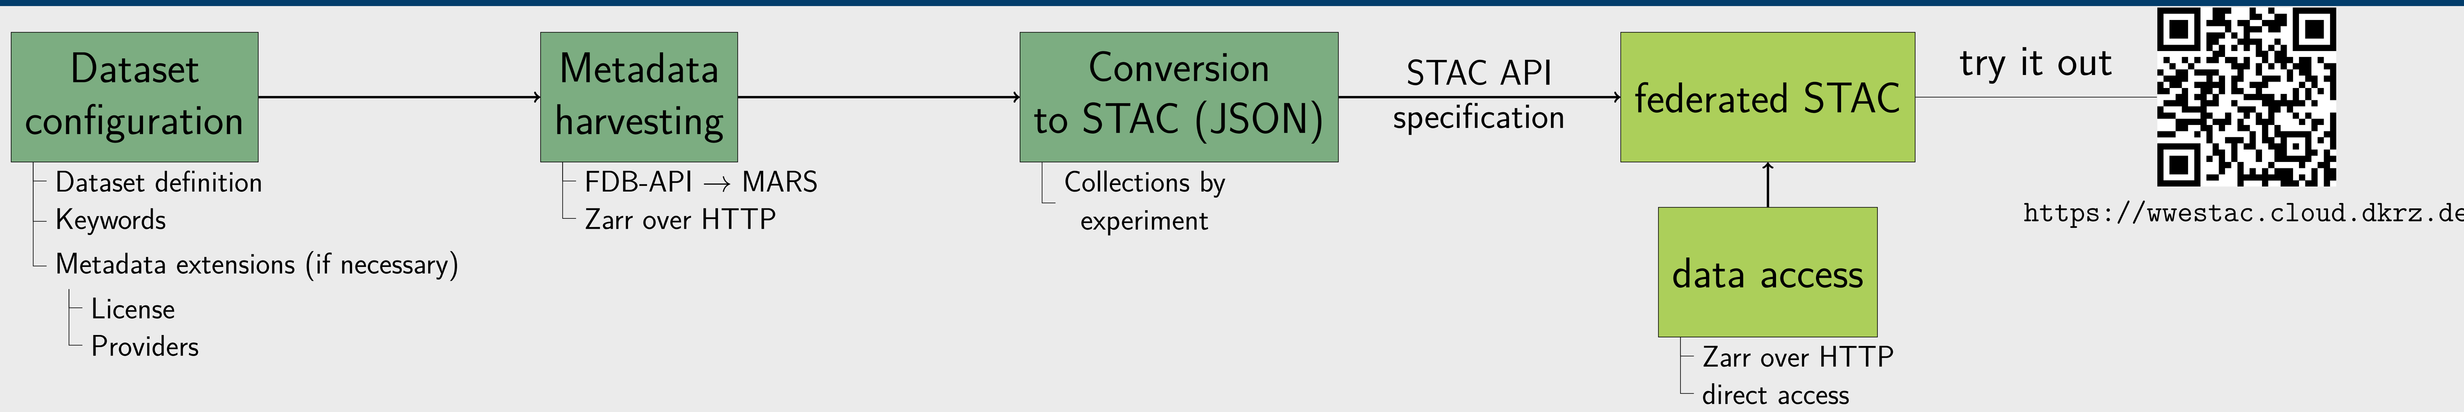

Examples: Easier Data Workflows for User

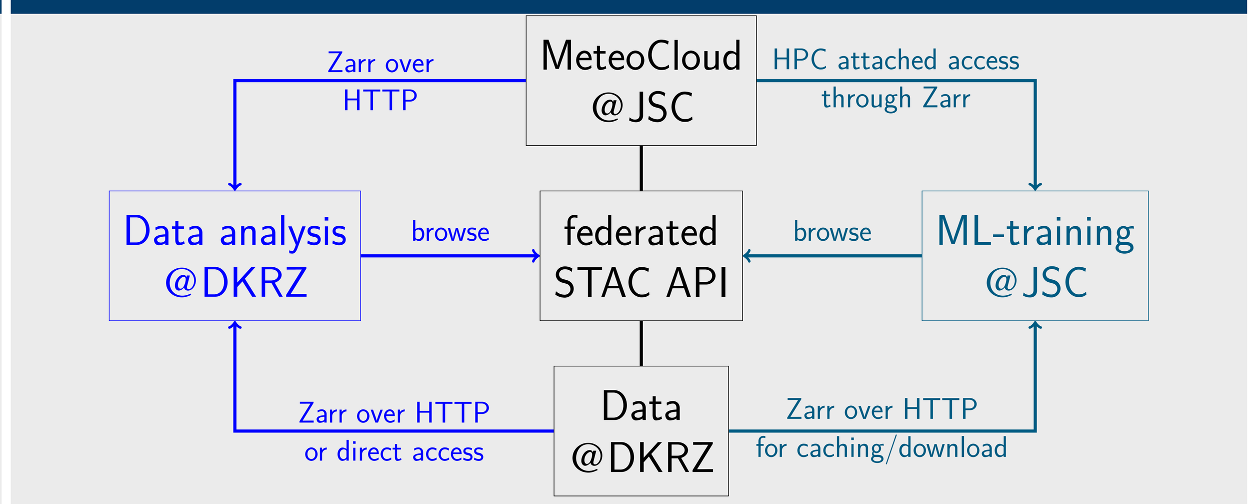

Summary

- Jülich MeteoCloud is already a large local repository
  - Jülich MeteoCloud on ExaStore now hosts ML-training ready datasets
  - Prototypical Zarr over HTTP workflows to make ESM data more accesible
- STAC improves the findability of data in the Jülich MeteoCloud (FAIR)
  - Creation of STAC entries from datasets
  - Shared catalogs improve discoverability also across data centers
- Extending federated data repository within WarmWorld

Outlook

- Extension of Jülich MeteoCloud with more datasets and storage
- Main challenge and performance risk: Inodes
  - Zarr3 with sharding for storage
  - Virtual Zarr implementations like kerchunk[4] or Cloudify[5]
- Service infrastructure to enhance Jülich MeteoCloud
  - Improved data access services
    - Deploy Zarr-FDB interface[6] as web service to open full datasets
    - Hosting of a *polytope*[7] service

  e.g. Automatic data movement between JUST6 and ExaStore as well as ExaFlash
- Towards ML-training metadata standards → Croissant[8]

References

[1] ECMWF. *anemoi-datasets documentation*. Apr. 21, 2026. URL: `https://anemoi.readthedocs.io/projects/datasets/`.

[2] Simon D. Smart, Tiago Quintino, and Baudouin Raoult. "A High-Performance Distributed Object-Store for Exascale Numerical Weather Prediction and Climate". In: *Proceedings of the Platform for Advanced Scientific Computing Conference*. PASC '19. Zurich, Switzerland: Association for Computing Machinery, 2019. ISBN: 9781450367707.

[3] *STAC: SpatioTemporal Asset Catalogs*. Apr. 22, 2026. URL: `https://stacspec.org/`.

[4] Martin Durant. *Kerchunk documentation*. Apr. 24, 2026. URL: `https://fsspec.github.io/kerchunk/`.

[5] Fabian Wachsmann et al. *Technical Overview of Cloudify - An Improved Emulator of Cloud-Optimized Earth System Model Output*. 2025.

[6] ECMWF. *Repository: Fields DataBase*. Apr. 20, 2026. URL: `https://github.com/ecmwf/fdb`.

[7] Mathilde Leuridan et al. "Polytope: an algorithm for efficient feature extraction on hypercubes". In: *Journal of Big Data* 12.1 (2025), p. 243. ISSN: 2196-1115.

[8] Mubashara Akhtar et al. "Croissant: A Metadata Format for ML-Ready Datasets". In: *Advances in Neural Information Processing Systems*. Ed. by A. Globerson et al. Vol. 37. Curran Associates, Inc., 2024, pp. 82133–82148.

Acknowledgements

We would like to thank our colleagues and project partners at ECMWF, especially Emanuele Danovaro, Tobias Kremer and Kai Kratz, for their discussions and support regarding our FDB setup. We would also like to thank Ilaria Luise for organising the distribution of data within the WeatherGenerator project and for facilitating discussions on the relevant metadata. Our thanks also go to Matthias Riße and Gebhard Günther from ICE-4 at FZJ for their work on data management and usage within MeteoCloud, and to Sabine Schröder (JSC) for her contributions to discussions on FAIRness and web services.

The authors would like to thank the German Federal Ministry of Research, Technology and Space for the funding of this work in the scope of WarmWorld: Easier (project identification number: 01LK2204D).

With funding from the: Federal Ministry of Research, Technology and Space

Sharing is encouraged

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