RDM for SeisSol using the Geo-INQUIRE SDL

2026-06-26

Readiness of HPC Extreme-scaling Applications (3rd Edition)

ISC HPC 2026 Workshop

iris.christadler@lmu.de, mathilde.marchandon@lmu.de, alice.gabriel@ucsd.edu

Geo-INQUIRE *is funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.*

SeisSol in a Nutshell

- **open-source** community code for **physic-based 3D dynamic rupture simulations of Earthquakes**
- uses **Discontinuous Galerkin discretization** and **local time-stepping on unstructured adaptive tetrahedral meshes.**
- Scalable performance at Petascale has been demonstrated **up to several thousands of nodes**
- Joint development by **LMU (Alice Gabriel, "Computational Seismology")** and **TUM (Michael Bader, "Hardware-aware algorithms and software for HPC")**

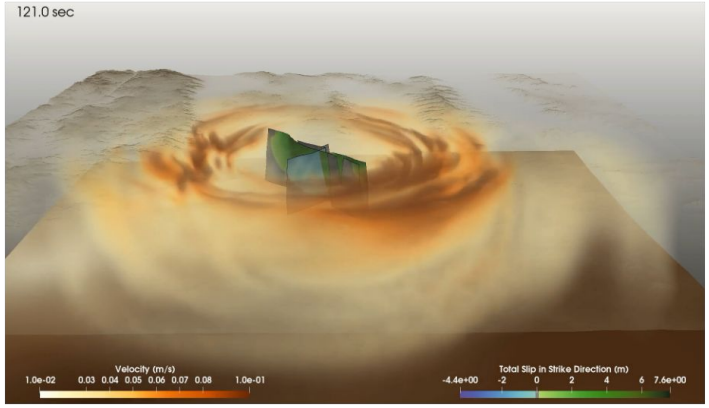

Propagation of seismic waves (2019 Ridgecrest earthquakes). Visualization of 15 TB of 3D volumetric data on unstructured tetrahedral meshes on Frontera. Taufiqurrahman et al., Nature, 2023
Abrams et al., SC23, Visualization showcase

SeisSol on Github: https://github.com/SeisSol/SeisSol/ doi.org/10.5281/zenodo.4672483

SeisSol Documentation: https://seissol.readthedocs.io/

SeisSol Training: https://github.com/SeisSol/Training
+video introductions online

Based on Slides from Alice-Agnes Gabriel

References:
Breuer et al.,ISC14, Heinecke et al.,SC14,
Breuer et al.,IEEE16, Heinecke et al.,SC16, Rettenberger et al., EASC16
Uphoff & Bader, HPCS'16, Uphoff et al., SC17
Wolf et al., ICCS'20, Uphoff & Bader, TOMS'20
Dorozhinskii & Bader, HPC Asia'21, Gabriel et. al., SC25

From “hero runs” to everyday examples

5 Hz Turkey simulation (larger area)

- required 9h on full (~8200 nodes) of Frontera,
- and a large mesh with 685 millions cell,
- + polynomial order 5 basis functions

Everyday Examples:Typical model size:

- meshes of ~10M elements,
- 200m spatial sampling of fault geometry,
- 200m finest sampling of topography,
- seismic wavefield up to 2 Hz in fault vicinity
- Comp. costs ~10,000 CPUh

Increases in compute power now allow for **many calibration runs, ”catalog generation”** and even **uncertainty analysis** with UM-Bridge > Increasingly data-rich

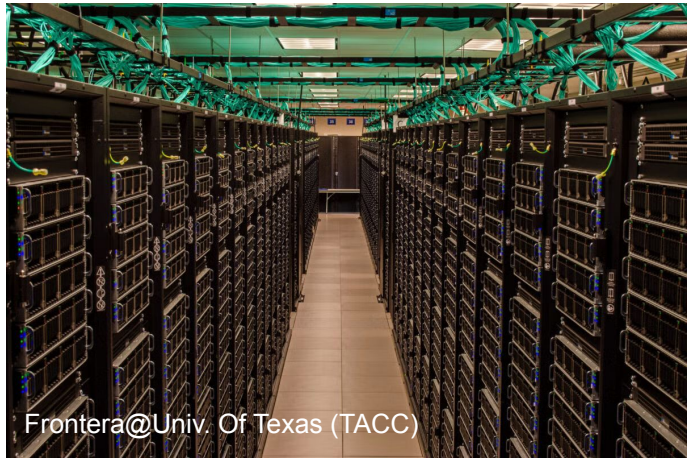
Frontera@Univ. Of Texas (TACC)


SuperMUC@LRZ

Based on Slides from Alice-Agnes Gabriel

Take-home messages

- **SeisSol is a “joint (ad)venture” of seismologists (LMU) and “hardware/numerics guys” (TUM)**
- **While SeisSol can be used for hero runs, more computing time is spent in “everyday” scenarios**
- **Increases in machine size and capability now enable bigger parameter studies leading to an significant increase in data**

- **This calls for unified workflows to store and access SeisSol data (input files, meshes, job scripts, outputs and data products)**
- **The Alto-Tiberina Earthquake Catalog (a DT-Geo “digital twin component” with a rapid response workflow) used the Geo-INQUIRE Simulation Data Lake (SDL)**

Usual approach takes weeks

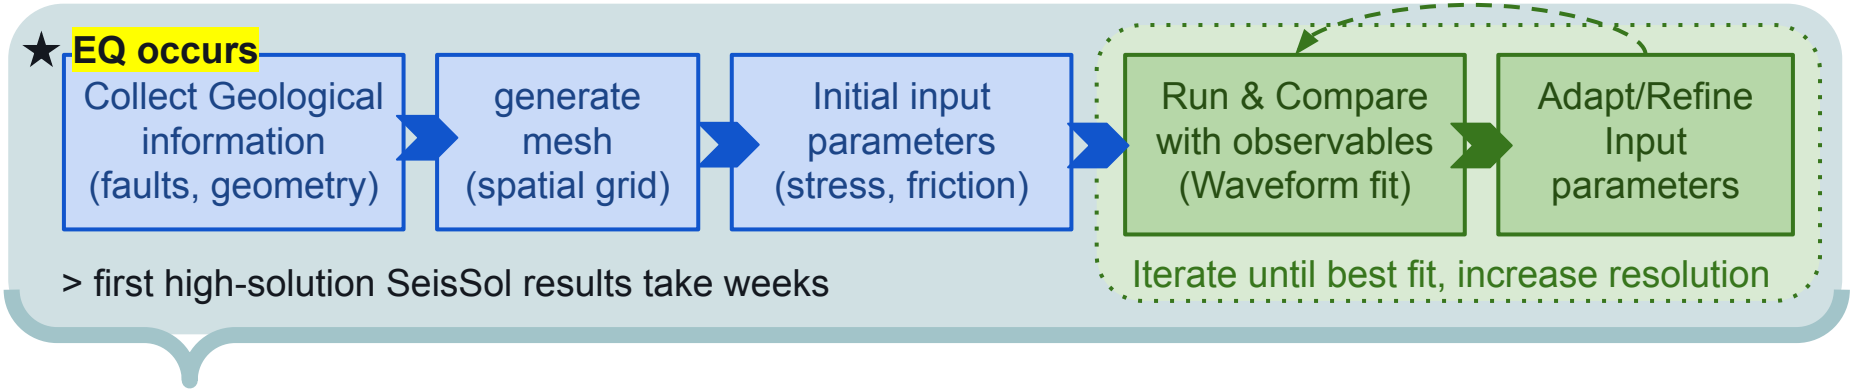

Solution: Pre-computed catalog for rapid response

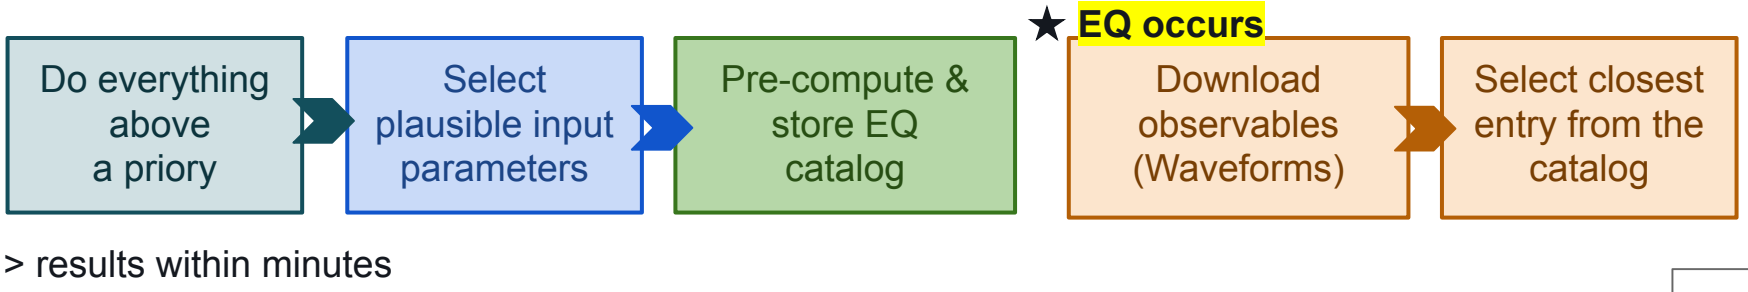

DT-Geo Demonstrator: Alto-Tiberina Catalog

Observation data & physics

-Fault geometry
-Stress orientation
-Frictional properties
-Medium properties

Multi-data constraints from Near-Fault Observatory **TABOO**

NEAR-FAULT OBSERVATORIES

Access data available within **min to hours** after the eventa

Moderate to **large event** occurrence

Simulation Workflow

SeisSol

3D dynamic rupture & seismic wave propagation

Varying nucleation location and values of unconstrained parameters

Catalog of **physic-based rupture scenarios…**

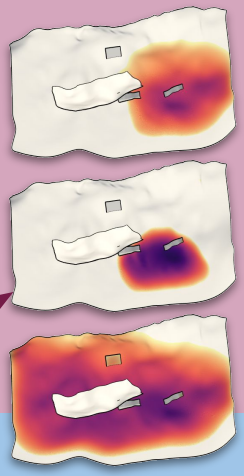

… with associated **synthetic data** (ground motion, static displacement, etc)...

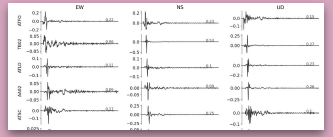

… and associated **shake maps**

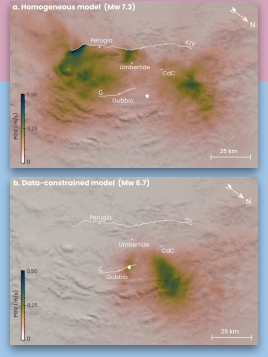

Search the catalog & output the best fitting scenario(s) (by comparing observations and synthetics)

FAIR Data for follow-on research

Alto-Tiberina Catalog: Catalog-Creation Workflow

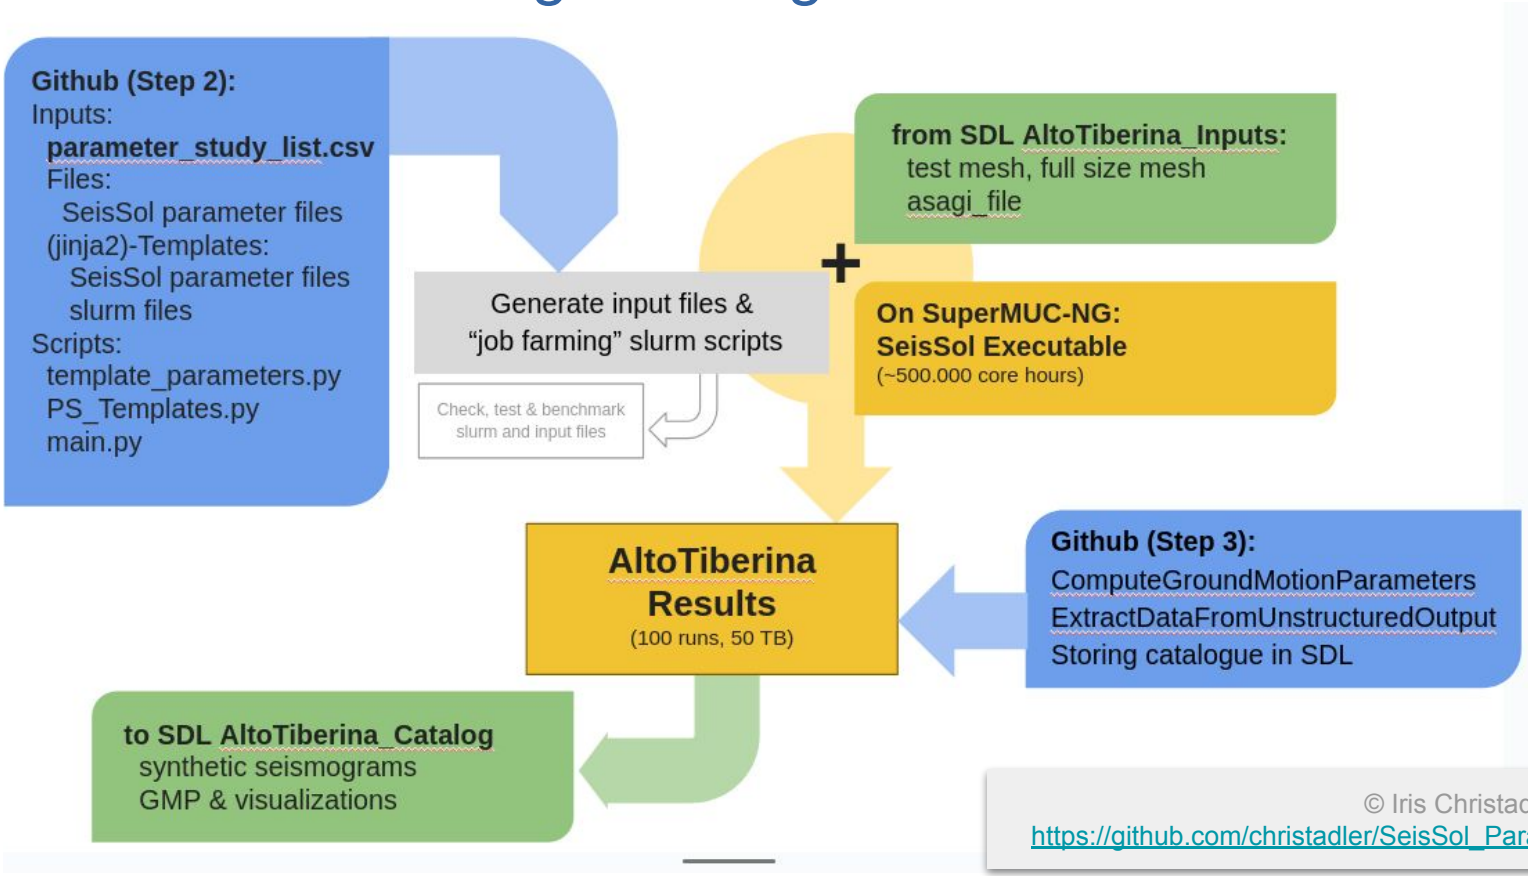

Alto-Tiberina Catalog: Rapid-Response Workflow

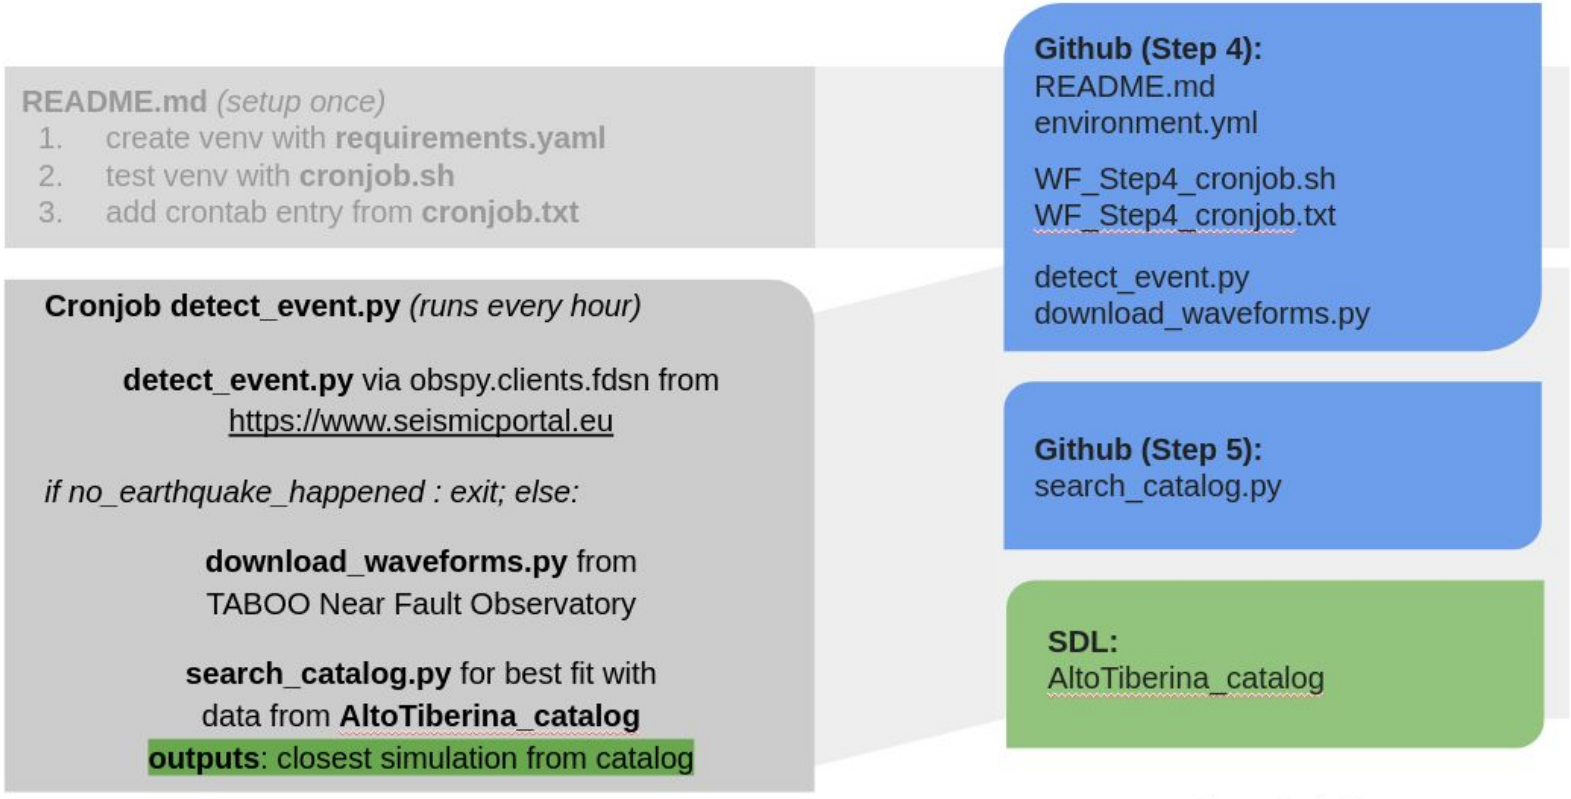

Geo-INQUIRE Simulation Data Lake

- **Data portal for simulated data** (inputs, results & data products)
- **FAIR:** Findable*Accessible*Interoperable*Reusable
- **Easy & fast access:** Upload GB in minutes
- **Metadata &DOI**, long-term storage
- Webinterface, CLI, Python, API
- Part of CINECA's AI Factory

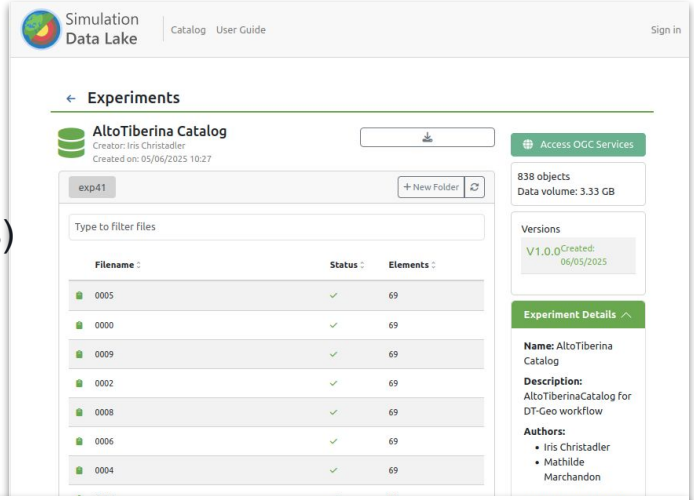

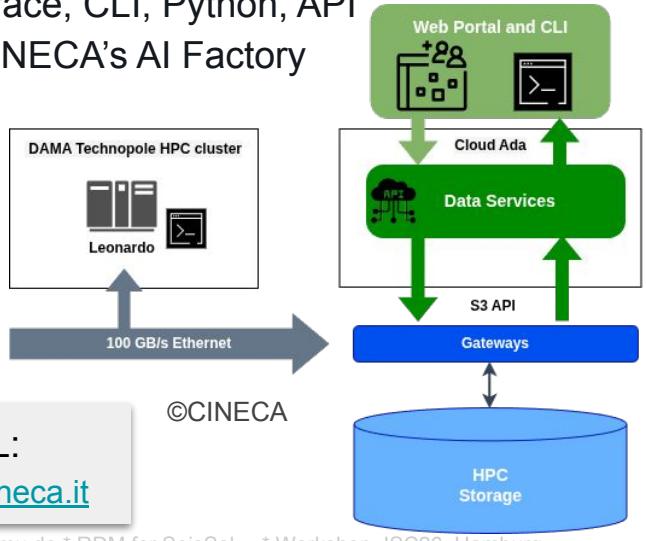

©CINECA

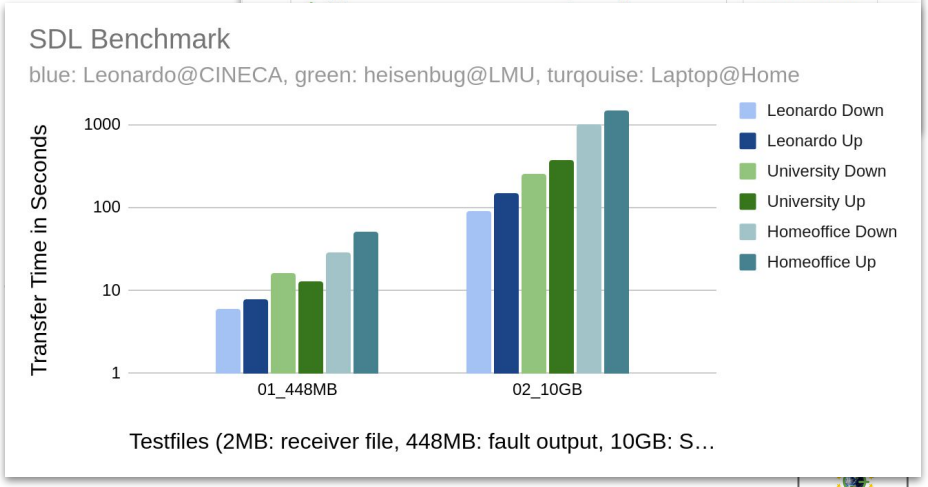

Access to SDL:
https://sdl.hpc.cineca.it

NFDI4Earth LUC Proposal

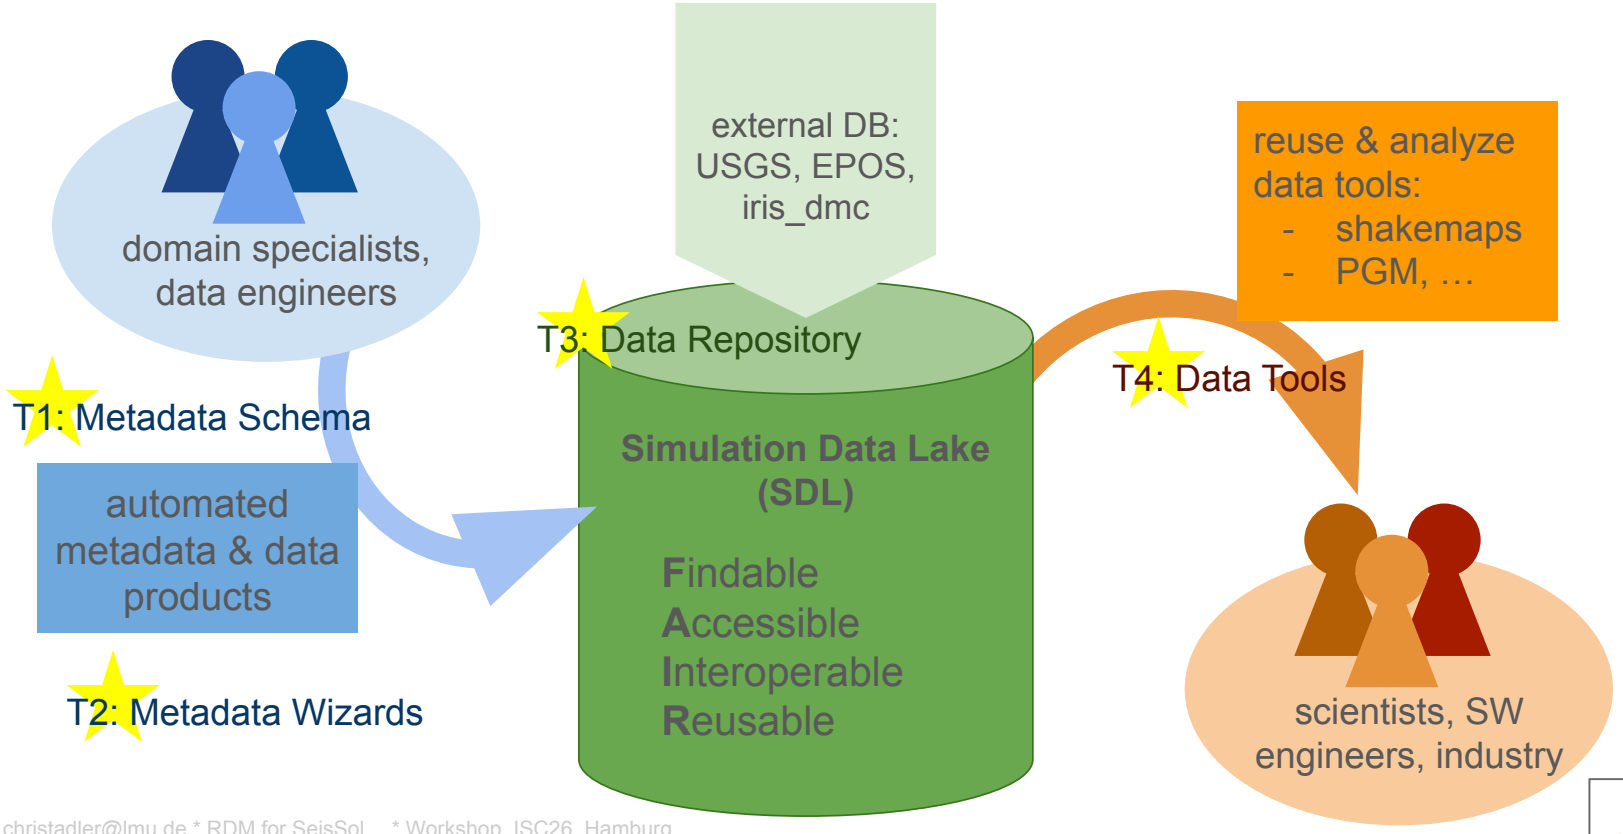

Take-home messages

- **SeisSol is a “joint (ad)venture” of seismologists (LMU) and “hardware/numerics guys” (TUM)**
- **While SeisSol can be used for hero runs, more computing time is spent in “everyday” scenarios**
- **Increases in machine size and capability now enable bigger parameter studies**
  Leading to an significant increase in data

- **This calls for unified workflows to store and access SeisSol data**
  (input files, meshes, job scripts, outputs and data products)
- **The Alto-Tiberina Earthquake Catalog (a DT-Geo “digital twin component” with a rapid response workflow) used the Geo-INQUIRE Simulation Data Lake (SDL)**

- **We are planning to provide an improved metadata schema and metadata “wizards”**
  to regularly store SeisSol outputs in a data lake

Thank you for your attention!

Geo-INQUIRE is funded by the European Commission under project number 101058518 within the HORIZON-INFRA-2021-SERV-01 call.