

Use case:

Processing the LOFAR Two-Metre Sky Survey (LoTSS)

- **LOFAR** radio interferometric data is large (~ 1 TB / hour)
- complex interplay of data calibration and sky brightness reconstruction („imaging“)
 - mixture of high-throughput and high-performance computing

Special requirements:

- data are stored on tapes (long-term archive)
- very large data volumes → having many copies unfeasible
- current software design requires powerful (single) nodes and sufficient scratch space ($\sim 5 \dots 20$ TB)



Use case:

Processing the LOFAR Two-Metre Sky Survey (LoTSS)



LOFAR LTA

dCache instance

davix-cp

GRID-certificate +
BEARER_TOKEN



condor_submit



c4p-login.gridka.de

**KIT compute
resources**

access token
renewal



Karlsruher Institut für Technologie

KIT dCache



Storage4PUNCH



Thüringer Landessternwarte
Tautenburg





Use case:

Processing the LOFAR Two-Metre Sky Survey (LoTSS)

Experiences using C4P and S4P:

- high demand for computing resources, large transfer time
 - nodes used so far (KIT) sufficient for processing single fields in a reasonable amount of time
 - storage size also sufficient for intermediate storage

What (in principle) prevents us to go for (semi-) automated bulk processing:

- manual token renewal prevents automation of the process in a long run
- independent service for data transfer to the compute nodes and back supporting standard protocols, e.g. gsiftp
- proprietary data if observed recently → user/group privileges management in S4P?



Thüringer Landessternwarte
Tautenburg

Radio Sky Image from LOFAR

