

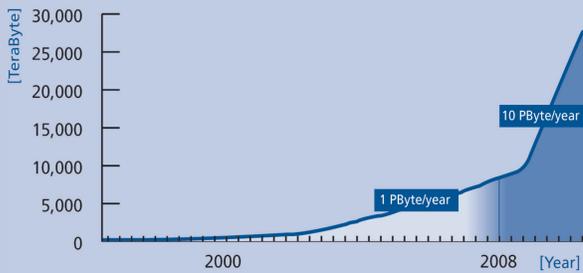
ICSU World Data Center Climate at DKRZ

DKRZ - DATA INTENSIVE CLIMATE SCIENCE



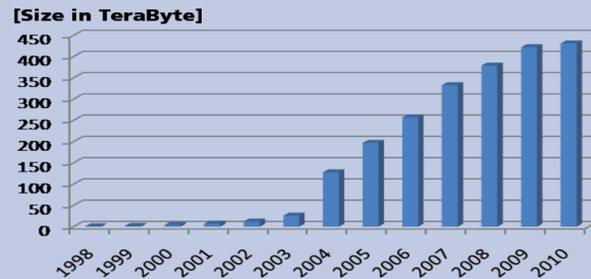
The World Data Center for Climate (WDCC) at DKRZ, that is a member of the ICSU World Data System (WDS), provides a long-term archive with approximately 500 TeraByte fully documented climate data. Users are provided with a flexible search interface and the possibility for field based data access and server side data processing (sub-setting, format conversion).

DKRZ archive over time



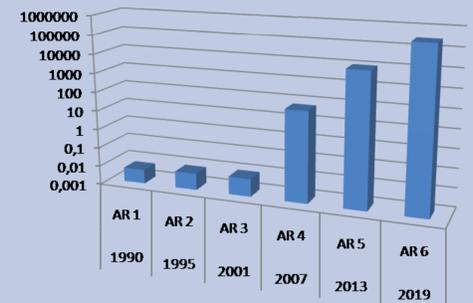
Exponential data growth at DKRZ. The DKRZ tape archive is designed for a growth of 10 PetaByte/year.

WDCC size over time



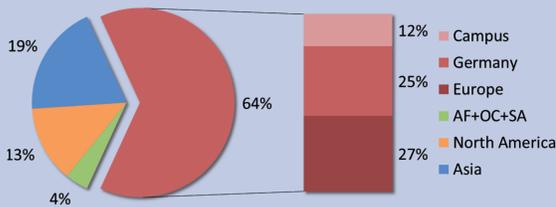
WDCC has grown by considerably over the last years and provides now access to more than 500 TeraByte of highly granular stored data. The WDCC is designed for a growth of 1 PetaByte /year.

IPCC data volume



Data volumes from IPCC have increased by a factor of 30 by each report. For the current AR5 report a volume of 10 PetaByte will be produced. For the next report in 2019 a volume of 300 PetaByte is expected.

WDCC: Users



Data is disseminated to users worldwide. Only a minority of users is located on the campus.

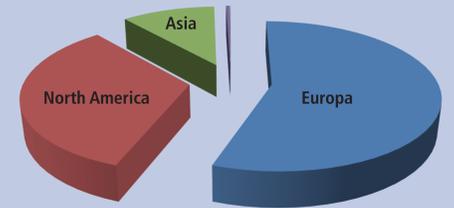
Challenges

Climate model simulations produce high volume multidimensional data products not optimised for direct end-user access.

Data and descriptive metadata have to be rearranged and made accessible supporting various usage scenarios and access patterns.

Thus the WDCC is based on a flexible metadata model as well as a generic data container system, to be able to efficiently serve data products ranging from few kiloBytes to hundreds of Gigabytes.

WDCC: Downloads



Downloads from Asia and North America represent a large part of the load.

Blades and disks

- 1.2 PetaByte storage system of DDN
- 10 Dell M605 Blades
- two AMD Opteron 4-Core 2.9 GHz
- ~32 or 64 Giabyte main memory
- 1 Dell M605 blade with two AMD Opteron

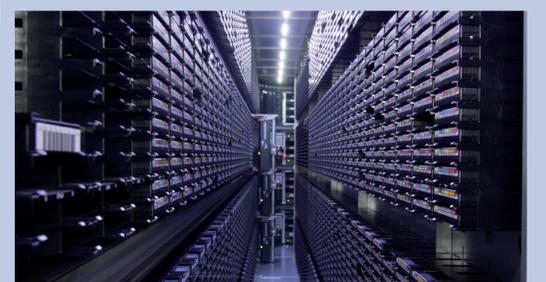


WDCC architecture and hardware

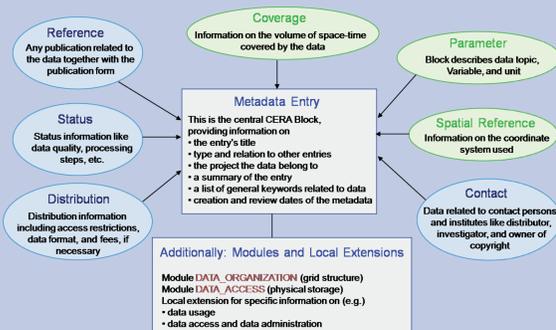
The WDCC is based on a complex metadata model (CERA), which is managed by an Oracle database system. Data is aggregated in container files, which are managed by so called LOB (large object) servers. These servers handle the staging of containers and (in cooperation with the CERA databases) support fine grained data access patterns.

HPSS - High performance storage system

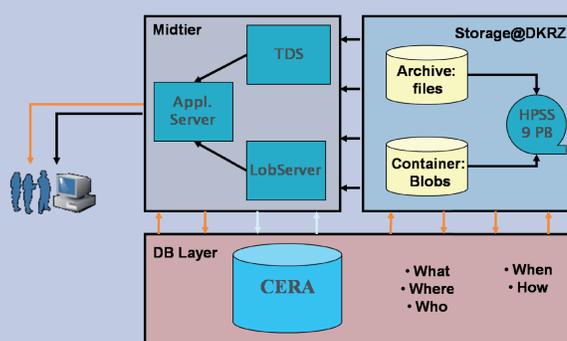
- 7 automated Oracle/StorageTek SL8500 tape libraries with more than 67 000 slots for magnetic tape cassettes
- 8 robots per library
- 88 tape drives
- more than 100 PetaByte total capacity
- bidirectional bandwidth of 5 GigaByte/s (peak)



CERA



Data access



Container

