ICOS Carbon Portal Usecase1: Footprint Tool

This is a short description of the work- and dataflow, storage and computation requirements of the usecase and its implementation with EGI and EUDAT services. It only serves as a basis for discussions between EGI, EUDAT and ICOS on the implementation of the usecase. The strategy is still open for discussion and the description is not yet complete. **More information will be added as required during the discussion.**

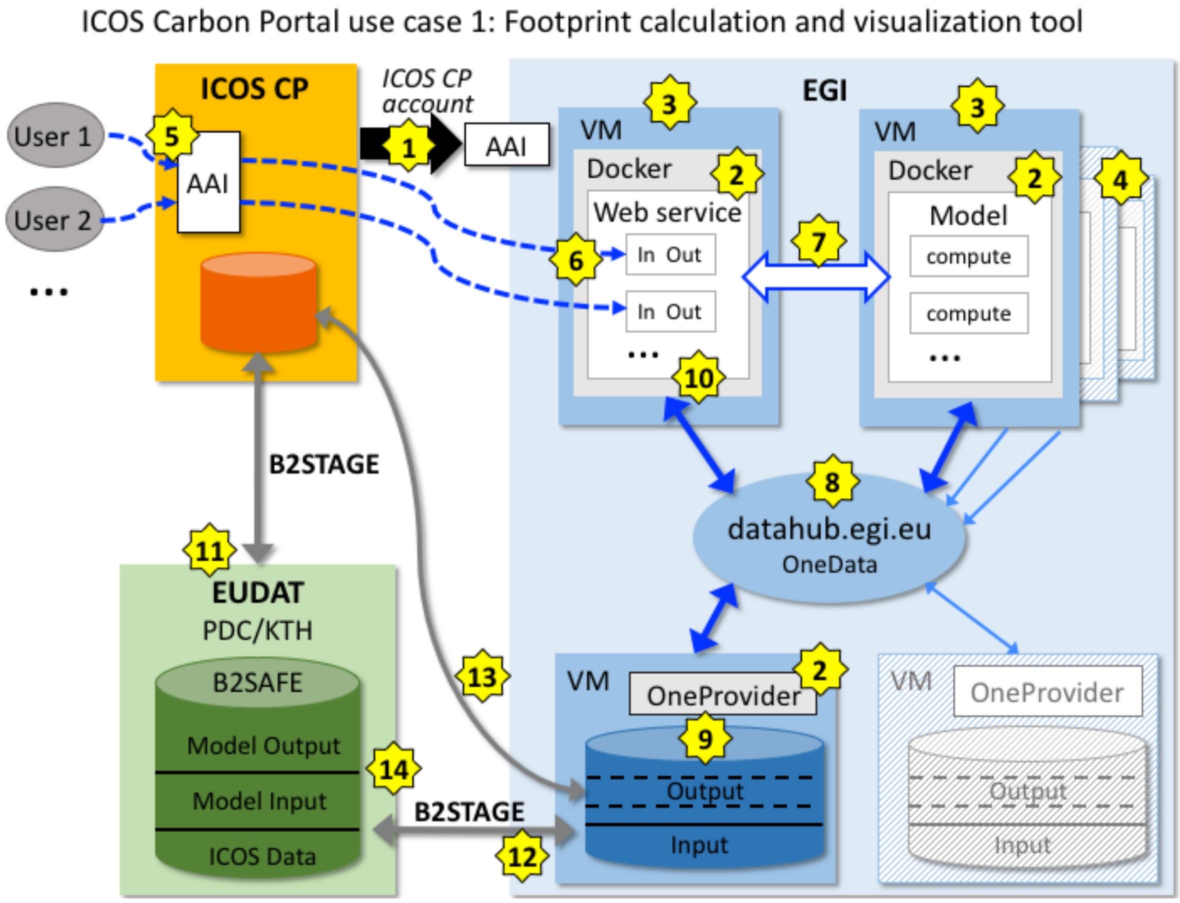

General work- and dataflow:

(Numbers correspond to those in figure above)

(Alternatives in italic letters

(Open questions highlighted)

Workflow for setting up and providing the footprint tool service:
(thick black arrow)

1. ICOS CP instantiates several VMs in the EGI FedCloud to host parts of the usecase.

1a. *A robot certificate associated to ICOS CP is used for the authentication and authorization in the EGI FedCloud.*

2. ICOS CP runs containerized versions of the services on these VMs:

- 1 VM hosting the web interface for request of model runs and visualization of the results.
- 1 VM hosting the transport model computations. (More VMs might be needed depending on the demand.)
- 1 VM hosting OneProvider with large block storage attached to store input and output data of the computations, which will be accessed by the VMs for visualization and computation. (More VMs might be needed depending on the demand.)

3. The VMs are running continuously and have constant IP addresses as long as they exists.

3a. *In case a permanent IP address would be available, this IP address is registered in the DNS of ICOS CP and the service is hosted on a subdomain of icos-ip.eu.*

3b. Otherwise a dynamic DNS service could be used. In that case, the ICOS CP DNS record for the service (e.g. stilt.icos-cp.eu) will not point directly at an IP address, but redirect to another domain name, which will be resolved by the dynamic DNS provider. The VM will then have to have the dynamic dns client installed, which will update the dynamic DNS record on VM startup.

4. In case several VMs are required to provide the service, load balancing is needed, e.g. docker swarm.

Workflow for a user of the footprint tool service:
(blue dashed arrows)

5. The user accesses the service at ICOS CP. Authentication and authorization of the user is handled at ICOS CP (HTTP cookie-based authentication or login). ICOS CP takes care of all interactions with the EGI FedCloud on behalf of the user.

6. The user is directed to the web service running in the docker on one of the VMs.

Dataflow inside the footprint tool:
(blue arrows)

7. The web service provides the user interface to start model runs and visualize results. Input parameters are selected on the web interface and passed to the script that starts the model run. The web service will serve several users in parallel and launch separate specific model runs.

7a. Model computations are hosted in a separate VM. As the model runs only on a single CPU, parallel model runs are possible. Communication between the applications would be done using HTTP.

7b. *Web interface and model computations are combined in one single VM.*

8. Model and visualization application both access input and output files on a virtual file system managed through the EGI datahub.

9. Input and output files are stored on block storage attached to a VM hosting OneProvider. New model output is merged with already available model output.

10. Model output is displayed on the web interface and the user can also download the results to her/his local computer, the web interface accesses input and output files on a virtual file system managed through the EGI datahub.

Data storage and transfer:
(grey arrows)

11. Quasi-static input data (incl. metadata) are stored on the ICOS CP server and for longterm storage in B2SAFE (at PDC/KTH), data transfer using B2STAGE.

12. A copy of the input data is transfered using B2STAGE from B2SAFE to the storage attached *to the VM hosting OneProvider* or directly to the local storage of each VM. Regular updates needed when new datasets become available.

13. Datasets not stored in B2SAFE, eg. intermediate versions of model output, are directly transferred from the ICOS CP server *to the VM hosting OneProvider* or directly to the local storage of each VM.

14. Model output data (incl. metadata) are transfered regularily to B2SAFE using B2STAGE and archived in B2SAFE for longterm storage.

The strategy to attach a PID (or DOI) to the model output (and user-specified request) is not yet included.

Workflow inside the footprint model:

Scenario 1: Visualization for existing ICOS sites
- Precomputed particle location files exist
- Footprints and time series exist for most sites and time ranges and emission types
- Check availability of footprints and time series files (if not available -> Scenario 2)
- Display results, very fast response time of web interface required

Scenario 2: On-demand calculation for new sites or new time periods
- No precomputed particle location files exist
- Full STILT run required (compute particle location, footprint, time series)
- Computations may require several hours
- Display results

Detailed workflow:
- Select time range and station (name-id or latitude/longitude)
- Check availability of aggregated footprints (netCDF-files) and of time series (csv-file) for full time range and all emission types
  - For each date (year, month, day, hour):
    - If footprint or particle location file not available
      - Compute particle location file if this is not yet existing (new station and/or new time period)
      - Compute footprint for this date and store (in netCDF-file)
    - Compute concentration for this date based on particle location and emissions, append result to csv-file

-> Display time series and animation of aggregated footprints in web service

Implementation of the visualization of footprints and time series at ICOS Carbon Portal:
https://data.icos-cp.eu/stilt/

Storage requirements:

- Input datasets (quasi-static, updated every 6-12 months)
  - Emissions (EDGAR, VPRM and more)
  - Initial and boundary concentration data
  - Meteorology
  - Particle location files (precomputed for many stations and dates)
  - Observation time series

Input dataset		no of files per year	File size	Storage per year	Total storage 10 years	Type
Emissions	EDGARv4.1	3	6.3 GB	19 GB	0.19 TB	netCDF
	EDGARv4.3	250	1-10 MB	< 1 GB	0.01 TB	netCDF
	VPRM	10	1-6MB	< 1 GB	0.01 TB	netCDF
	others	3-10	20-100 MB	< 1 GB	0.01 TB	netCDF
Boundary		5	1-10 GB	20 GB	0.20 TB	netCDF
Meteo		12	17 GB	204 GB	2 TB	arl (binary)
Particle location per station	1-3 hourly	2920-8760	0.5-5 MB	~ 20 GB	~ 20 TB	R-object
all available	90 stations	~ 400000	0.5-5 MB	~ 2 TB		

- Output datasets
  (model runs access output data from previous runs and eventually add new files)
  - Aggregated Footprints
  - Concentration time series
  - Particle location files for new sites (produced in a full STILT run)

Output dataset		no of files per year	File size	Storage per year	Total storage 10 years	Type
Footprints per station	1-3 hourly	2920-8760	40 kB	< 300 MB	~ 150 GB	netCDF
all available	90 stations	~ 400000	40 kB	~ 15 GB		
user requests	new sites			~ 20 GB	~ 200 GB	
Time series per station	1-3 hourly	1	1-10 MB	~ 10 MB	~ 10 GB	csv (ascii)
all available	90 stations	90	1-10 MB	< 1 GB		
user requests	new sites			10 GB	~ 100 GB	
Particle location per station	1-3 hourly	2920-8760	0.5-5 MB	~ 20 GB	------	R-object
user requests	new sites			~ 2 TB	< 20TB *	

*User requests will initiate computation of additial footprints and time series for approx. 100 new sites but for only 1-2 years.

Computational requirements:

Test on a linux cluster at MPI for Biogeochemistry in Jena and VM in EGI FedCloud
(*will add specification here)*

Full STILT run:
3 GB memory per job
670 CPU seconds per footprint
1700 CPU hours per station per year
(more detailed information on cpu time etc. will be added)

Model runs for individual stations and users requests are separate jobs therfore paralell processing is possible and required for better performance.