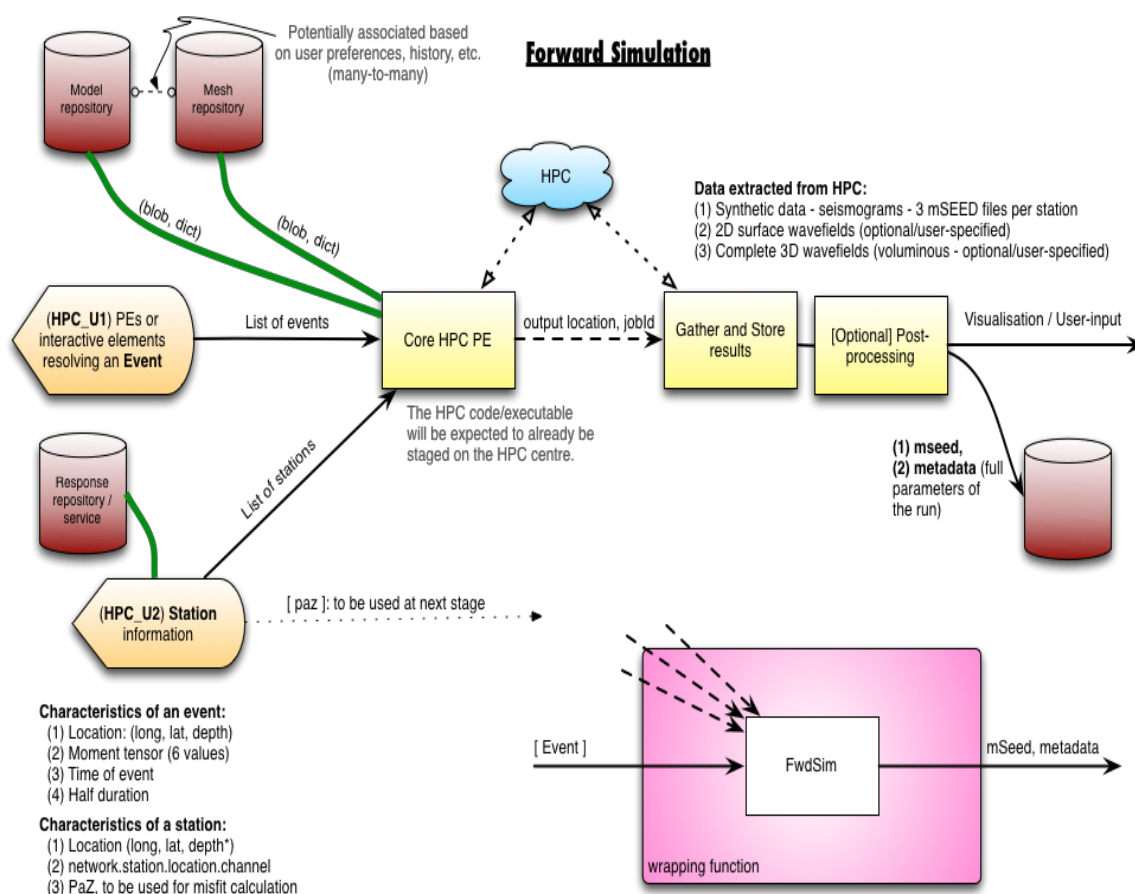


Generation of synthetics - data storage and discovery requirements



Driver Requirement 1

The above diagram is part of one of VERCE's use-cases and its purpose is to generate synthetic data according to a mathematical model. Such models are meant to be executed on HPCs (currently CINECA PLX and LRZ SuperMuc). The yellow and orange boxes represent processing elements (PEs), or workflow subgraphs, typically to be executed at locally-managed/private computing resources.

In order for such execution to take place, PEs need to be provided with a list of stations (seismic stations and their locations) and a list of seismic events. Basic lists of characteristics upon which such entities should be searchable or 'discoverable' are shown on the graph. Given this information we can retrieve the corresponding seismograms - series recordings of accelerations or misplacement. Depending on the number of stations (geographical region) and time constraints of events chosen, data volume can range from a few hundreds of MBs to a few tens of GBs - these estimations are currently very crude as we don't have a fully working system yet to test under realistic conditions).

Current practice: some of this information is currently served up by online/web services and some of it might reside, in the form of files, in data storage resources of VERCE participants. We would potentially like to be able to:

(a) normalise the discovery procedure regardless of location, by providing an internal VERCE service that interfaces between various external or internal services and the workflow/user.

(b) as these form input parameters for a procedure to take place on HPC resources, we would like such discovery to be followed by a transfer to the appropriate resource directly, before a job can commence (which would require some integration between the discovery mechanism employed, storage or mediator services and HPC middleware).

Driver Requirement 2

After the experiment has run successfully on the HPC (past the 'Core HPC PE' of the diagram), a series of synthetic seismograms have been generated on the HPC. Depending on resolution and parameters (whether they produce 2D or 3D wavefields), such synthetics can range from a few GBs to many hundreds of GBs.

The synthetics then need to be staged out of the HPC and into some semi-permanent storage from where they can be streamed into the remaining of the workflow for post-processing. This, at the moment, is being orchestrated from appropriate processing elements using gridFTP, as appropriate.

At the end of post-processing, and potentially after approval by the scientist, the synthetics need to be stored in a more permanent way. The synthetics will have to carry (i.e. be annotated by) and be discoverable by all the relevant metadata used during processing - lists of stations and events, models, meshes, characteristics of pre- and post-processing PEs, etc., in the form of metadata. Such metadata will be used in the future to cross-check the validity of the synthetics by re-running the same experiment under the same conditions and also for latter stages of the use-case, where they will be compared against real data in order to tweak the model towards a more realistic one.

Resource Discovery and Related Information

In VERCE we would like to 'hide' resource discovery from the end-user, where appropriate. During simulation runs, resource discovery should be handled by the Dispel/Workflow system, as this would be part of the optimisation procedure. For this optimisation to work, VERCE components should know of their surroundings, including access points for storage and computing facilities, load, disk space quotas (where appropriate), etc. Perhaps the GLUE specification would be a good starting point for making these requirements more explicit.

Current Technology Considerations

Storage

The main technology we are currently considering is IRODS in that it offers file-storage along with key,value metadata. Due to potentially large output files, we are also considering some IRODS solution over a DFS (e.g. HDFS/HBASE for archiving). The actual computing resources to be used for this purpose are still under discussion but, realistically, we would probably need to 'outsource' some of these storage, management and discovery requirements. Working within the EGI/EUDAT/PRACE synergy could potentially provide us with some automation and transparency regarding discovery and staging at PRACE sites, grids, etc.

A&A

This will be certificate-based. The core entities we anticipate to have in VERCE are Users and Institutes (consisting of users) and we would like to support collaboration between users and institutes.

Data discovery and storage also needs to be driven by the identity of the user and the resources allocated to the institute. For instance, users should be able to run experiments only discoverable by themselves, or shared across the VERCE infrastructure.