



# Staging Replicas for Computation and its Requirement

Stefan Zasada  
University College London



Date: 7<sup>th</sup> March 2012




## Motivation


- Many users have data stored on EUDAT infrastructure and need to perform analysis using a subset of this data  
→ e.g. simulations or data mining.
- The data needs to be moved from the EUDAT infrastructure to the resource doing the processing  
→ This may be a PRACE machine, an EGI cluster or a local machine.
- The data collection is likely to be large, and will be needed at the compute site for an extended period of time
- When the simulations/analysis has finished results need to be brought back to EUDAT

# Scenario 1: All simulations are run on the same PRACE system

Data set could range from a few TBs to 100-1000s of TBs, with 10k plus individual files



## Scenario 2: Multiple simulations are run on multiple PRACE systems





## Detailed workflow

1. Replicate data collection from master storage to storage close to HPC system(s).
2. Run simulation by copying data from storage to HPC workspace and start the simulation.
3. Copy results from the simulation back into the storage.
4. Replicate the results to slave archives also working on the same data collection.
5. Replicate the simulation results to the master.
6. Slaves notify the master about all extra copies.
7. The master creates or updates the PID records based on the notifications from steps 5 and 6.



# Requirements

- Users can specify which collections to replicate for a simulation.
- Users should be able to specify which EUDAT centers to use for the dynamic replication. Preferably close to the HPC system.
- Users can specify how long the data should be kept close to the HPC system.
- Data is moved from the EUDAT storage to the HPC workspace.
- Replicas across multiple HPC centers should be kept in sync once a simulation is run in one of the centers (step 4 in scenario 2).
- PID Server returns optimal URL (see PID service case description for details).
- Community Managers should be able to manage user permissions.
- Community Managers want to know whether the replicas are identical to the source (auditing).
- There is the need to control what user can do in terms of starting replications to and simulations on HPC systems, restrictions on how long user can keep data in the EUDAT storage...



## Summary

- Initial requirements specified by Stefan Zasada, Ali Haidar, Willem Elbers, Peter Wittenburg
- Requirements shared by VPH, ENES, EPOS & CLARINS

Questions?