

Ystia Orchestrator and HEAppE Middleware for Earthquake and Tsunami HPC workflow management

Jan Krenek¹, Laurent Ganne², Vaclav Svaton¹, Jan Martinovic¹, Katerina Slaninova¹, and Thierry Goubier³

¹*IT4Innovations, VSB – Technical University of Ostrava, Ostrava, Czech Republic*

²*ATOS, Grenoble, France*

³*CEA LIST, Palaiseau, France*

1 LEXIS Project

The LEXIS project (<https://lexis-project.eu>) is developing an advanced online engineering platform integrating HPC, Cloud and Big Data. The platform will leverage large-scale geographically distributed resources from existing HPC infrastructures, employ Big Data analytics solutions and augment them with Cloud services. The platform integrates the Ystia Orchestrator: a workflow designer and orchestrator developed by Atos and the HEAppE Middleware: HPC-as-a-Service concept implementation developed by IT4Innovations National Supercomputing Center. The poster will illustrate the architecture of the 'workflow & orchestration' part of the platform and its applicability for the LEXIS Earthquake and Tsunami pilot use case.

1.1 Ystia Orchestrator and HEAppE Middleware

The Ystia Orchestrator (<https://github.com/ystia/yorc>) is an open source TOSCA (Topology and Orchestration Specification for Cloud Applications, an OASIS consortium standard) orchestrator. It supports applications/jobs lifecycle management over hybrid infrastructures (HPC schedulers, Kubernetes, OpenStack, several public clouds), for which it provides an implementation of operations allowing to create/delete infrastructure resources (jobs, compute instances, block-devices, etc...) on demand. A front-end provides a UI and a REST API allowing to easily create applications and workflows, from an extensible catalog of TOSCA components.

Creating generalized AAI (Authentication and Authorization Infrastructure) across multiple service providers is always a challenge due to differing security policies and mechanisms. The HPC-as-a-Service concept offers high-level abstraction and simplification in the area of HPC computing for regular users and novice users with no HPC knowledge whatsoever. The HEAppE Middleware (<http://heappe.eu>) is IT4Innovations' in-house implementation of this concept. Thanks to its mapping mechanism between external user accounts and internal cluster accounts, it provides access to HPC capabilities with additional secure access layer. Computational jobs are managed by the external user accounts (standard users, not registered in a processing centre) via the HEAppE Middleware but the actual job submission to a processing queue (PBS/Slurm) is done under service/anonymous HPC cluster accounts. This approach greatly simplifies and accelerates the overall user registration process and access to HPC resources while respecting the processing center's internal security/ISO processes. From the security point of view the computational project admin is responsible for these service accounts and how they are used by the HEAppE, HEAppE instance itself has to be located within the center's network and these internal service accounts should not leave the premises of the center so the center's support staff always know how they are used.

2 Workflows and Orchestration

The Ystia Orchestrator is workflow driven: it doesn't contain any hard-coded lifecycle logic. This allows a user to fully customize applications behavior, and to provide any custom workflow that can be executed at runtime. The Ystia front-end can generate automatically installation/uninstallation workflows from the description of components in your applications and their relationships (declarative workflows), or workflows can be user-defined (imperative workflows). A workflow is a sequence of operations occurring on possibly different types of infrastructure, like for example running a job on a HPC cluster and then creating a VM in the cloud and deploying software on it. Policies (placement, failover, scalability) can also be specified to add automatically additional constraints to the workflow.

3 Future Work

To take advantage of the federation of resources from geographically distributed HPC infrastructures, the following innovations will be brought in LEXIS orchestration:

- dynamic orchestration: select the location where to allocate resources needed in a workflow dynamically during the workflow execution, according to placement policies,
- job failover: synchronize checkpointing data between locations, perform a failover to another location in case of failure,
- job scalability for urgent computing: submit a job on several locations in parallel, have the first submitted or run job cancel the others,
- deadlines for urgent computing: the ability to trigger changes or processing linked to world-time deadlines, such as cancel a job, or submit according to the time left before the deadline.