

Orchestration of a Forecasting Chain for Forest Fire Prevention using the LEXIS Cloud/HPC Platform

Mohamad Hayek
 Leibniz Supercomputing Centre (LRZ)
 Bavarian Academy of Sciences and
 Humanities
 Garching b. M., Germany
 ORCID: 0000-0001-8607-6329

Laurent Ganne
 ATOS
 Grenoble, France

Antonio Parodi, Andrea Parodi,
 Mirko D'Andrea
 CIMA Research Foundation
 Savona, Italy

Rubén J. García-Hernández, Stephan
 Hachinger, Johannes Munke
 Leibniz Supercomputing Centre (LRZ)
 Bavarian Academy of Sciences and
 Humanities
 Garching b. M., Germany

Martin Golasowski,
 Kateřina Slaninová
 IT4Innovations
 VŠB -
 Technical University of Ostrava
 Ostrava, Czech Republic

Frédéric Donnat
 Outpost24
 Antibes, France

Emanuele Danovaro
 European Centre for Medium-Range
 Weather Forecasts
 (ECMWF)
 Reading, United Kingdom

Olivier Terzo, Alberto Scionti,
 Giacomo Vitali
 Advanced Computing and
 Applications
 LINKS Foundation
 Turin, Italy

Jan Martinovič
 IT4Innovations,
 VŠB -
 Technical University of Ostrava
 Ostrava, Czech Republic

Abstract — We describe a first successful application of the “LEXIS” platform for advanced orchestration of complex simulation and data-analysis workflows in mixed Cloud/HPC environments. A workflow for forest-fire risk assessment based on the models WRF and RISICO was executed, using IaaS-Cloud resources at LRZ (DE) and national HPC resources at IT4I (CZ). Besides concept and results of the workflow, our contribution focuses on core systems of LEXIS: A “Distributed Data Infrastructure” using iRODS and EUDAT, and an Orchestration System with the novel ability to address HPC schedulers besides IaaS Clouds. The approach and open platform, currently federating IT4I, LRZ and ECMWF (UK), are aimed at broad application within the current project (“Large-scale EXecution for Industry and Society”, H2020 GA 825532) and beyond.

Keywords — Workflows, High-Performance Computing, Cloud Computing, Simulations, Orchestration, Data Management

I. INTRODUCTION

The LEXIS [1] project aims at building an advanced, user-friendly platform for the execution of demanding simulation and data-analysis workflows. It takes advantage of optimized orchestration and data management techniques, and state-of-the-art computing infrastructure. Initially, LEXIS supports three use cases from science and industry. These “Pilots” focus on demanding workflows related to (i) turbo-machinery and gearbox simulations, (ii) earthquake and tsunami data processing, and (iii) weather and climate models. In this paper, we discuss the LEXIS platform and a first successful execution of a workflow for forest-fire risk assessment.

II. DISTRIBUTED DATA INFRASTRUCTURE

Orchestration of workflows over geographically distributed Cloud and HPC systems (currently at IT4I/CZ and LRZ/DE), as offered by LEXIS, requires consistent data

management. The LEXIS Distributed Data Infrastructure (DDI) provides a unified file-system-like view on all data in the project, by leveraging the “Integrated Rule-Oriented Data System” (iRODS, [2]) middleware on top of the distributed back-end systems. iRODS was chosen because of its flexibility and suitability for integration with the European DATa (EUDAT, [3]) infrastructure. Our current iRODS federation consists of two “zones” (IT4I & LRZ). On top of iRODS, we deployed the EUDAT services B2HANDLE (persistent identifiers, [4]), B2SAFE (data safety and fast access by replication, [5]), and B2STAGE (data transfer to/from HPC and GridFTP, [6]).

REST APIs are developed in LEXIS as a front-end to the DDI, offering endpoints used by the Orchestration System and User Portal. A staging API provides methods to move data from/to Cloud and HPC resources. Metadata annotation APIs allow for a data management following FAIR [7] principles. The LEXIS platform is sketched in Figure 1. The DDI connects to the LEXIS Authentication and Authorization Infrastructure, providing access control via iRODS ACLs.

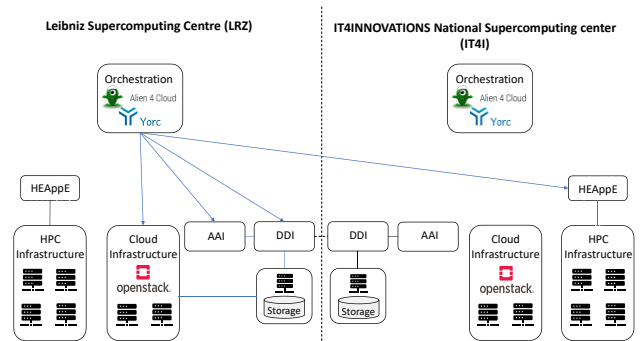


Figure 1: Distributed LEXIS infrastructure; blue arrows mark components addressed by the orchestrator in our WRF Workflow (cf. Sections IV-VI).

III. AUTHENTICATION AND AUTHORIZATION IN LEXIS

The LEXIS Authentication and Authorization Infrastructure (AAI), which all LEXIS services rely on, is addressed by the DDI via the OpenID Connect [8] mechanism. The core of the AAI is a central service running Keycloak [9] and providing a unified user login for the LEXIS platform.

Users will access LEXIS resources via a LEXIS Portal (already mentioned above), as a user-friendly entry point to the platform with its data and compute services. The user authenticates in Keycloak, and receives a JWT [10] token, which is forwarded to all relevant back-end systems, including the orchestrator and the DDI. The token encodes the clearances and access permissions of the user. To render the token-based access to all services secure, the token and its validity are verified against Keycloak on every exchange.

IV. LEXIS ORCHESTRATION AND WRF WORKFLOW

LEXIS helps users to easily automatize complex, hybrid Cloud/HPC workflows with a novel Orchestration System. Applications/Workflows are automatized by creating models using TOSCA [11] (Topology and Orchestration Specification for Cloud Applications). TOSCA describes an application workflow made of components, with their relationships, requirements, capabilities, and operations. The orchestrator is based on the Yorc solution [12] (back-end), and Alien4Cloud [13] (front-end). To launch and monitor workflow components on both Cloud and HPC infrastructure, an appropriate orchestrator plugin was developed. It addresses resources through the HEAppE HPC middleware [14]. The system allows for workflow control via REST API calls or UI. The LEXIS Orchestration System will be accessible via the LEXIS Portal. With few clicks, a user can choose a workflow, input data, and other parameters and execute the workflow.

One of the first applications automatized with the LEXIS platform is a workflow for forest fire prevention with the WRF (“Weather Research and Forecasting”, [15]) model as “HPC core”. This workflow starts with GFS data acquisition from NCEP servers and a WRF preprocessing system (WPS) run on the LRZ OpenStack [16] IaaS Cloud. WPS runs on the cloud as it requires limited, easy-to-use computational resources. Afterwards, a WRF HPC job is created and executed on IT4I’s Salomon Cluster using the WPS results. After this is completed, data are staged back to LRZ, where a forest-fire decision support framework (“RISICO” [17]) is executed on the LRZ cloud. The staging API is used to move the final results to the DDI. The workflow steps are depicted in Figure 2. The results are finally moved to further systems, e.g. for visualization, via the staging API and EUDAT B2STAGE. The total execution time of the workflow is 6 hours using 24 hours of forecasts data. Running the workflow manually requires at least 7 hours using exactly the same resources.

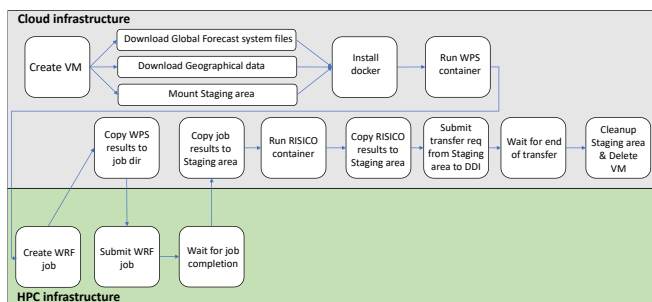


Figure 2: Forest fire prevention workflow as executed by the orchestrator.

V. DATA STAGING IN AN ORCHESTRATED WORKFLOW

Within the workflow shown above, the Orchestration System moves all relevant data between the HPC, Cloud and data infrastructures involved, by triggering asynchronous data staging jobs via the staging API of the DDI (Sections II, IV). The Orchestration System can then query the status of the jobs via another API endpoint. Additionally, the API provides an endpoint to delete data on different targets once the workflow is executed successfully. To implement asynchronous file transfer and deletion in the Staging API, a Celery [18] worker runs behind the Django [19] based frontend. The Celery jobs are managed by a RabbitMQ [20] queuing system.

VI. WORKFLOW RESULTS

The results of the workflow sketched above can conveniently be visualized with the Dewetra platform [21] of the Italian Department of Civil Protection. Dewetra renders various parameters relevant to forest fire prediction and prevention, and is ideally fed by scheduled (e.g. daily) runs of our workflow. The visualization includes Fireline Intensity and the Rate of Spread which is affected by winds. Figure 3 shows an example for Italy on the 17th of June, 2020. On that day, the Fireline Intensity was high over Basilicata, Calabria, northern Italy, and Sardinia. The predicted Rate of Spread peaked in southern Italy, parts of Sicily, and again in Sardinia.

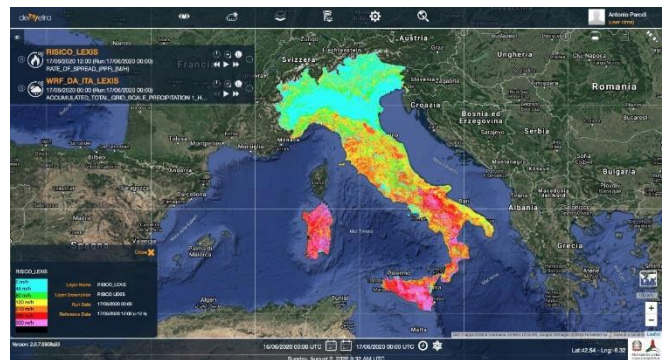


Figure 3: Workflow results visualized with the Dewetra platform.

VII. CONCLUSION AND FUTURE WORK

A Pilot workflow for forest fire risk assessment and the LEXIS platform to execute it was shown. The workflow relies on the RISICO framework and HPC weather simulations with WRF. The LEXIS platform runs this on large Cloud and HPC systems at two PRACE [22] Tier-0/1 sites, using an efficient Orchestration System, a Distributed Data Infrastructure, and a common Authentication and Authorization Infrastructure. A web portal is being developed as a user-friendly front end. The LEXIS platform shall support a large variety of workflows in the near future, so we will focus on benchmarking and optimization of the platform, and on co-development considering the demands of new users. We have also launched an Open Call to attract use-cases from life sciences, manufacturing, and other fields significant for industry and society.

ACKNOWLEDGMENT

This work and all contributing authors are (co-)funded by the EU’s Horizon 2020 research and innovation program (2014-2020) under grant agreement no. 825532 (Project LEXIS – “Large-scale EXecution for Industry and Society”).

REFERENCES

- [1] A. Scionti et al., „Cloud and Big-Data Convergent Architectures: the LEXIS Approach,“ in *Proceedings of the 13th International Conference on Complex, Intelligent, and Software Intensive Systems (CISIS-2019). Advances in Intelligent Systems and Computing*, vol. 993, Sydney, Australia, 2019.
- [2] H. Xu, T. Russell, J. Coposky, A. Rajasekar, R. Moore, A. de Torey, M. Wan, W. Shroeder und S.-Y. Chen, *iRODS Primer 2: Integrated Rule-Oriented Data System*, Williston, VT: Morgan & Claypool Publishers, 2017.
- [3] EUDAT Collaborative Data Infrastructure, "EUDAT," EUDAT Ltd, 2020. [Online]. Available: <https://www.eudat.eu>. [Accessed 27 Apr. 2020].
- [4] EUDAT Collaborative Data Infrastructure, "B2HANDLE - EUDAT," EUDAT Ltd, 2020. [Online]. Available: <https://www.eudat.eu/services/b2handle>. [Accessed 13 Apr. 2020].
- [5] EUDAT Collaborative Data Infrastructure, "B2SAFE - EUDAT," EUDAT Ltd, 2020. [Online]. Available: <https://www.eudat.eu/services/b2safe>. [Accessed 13 Apr. 2020].
- [6] EUDAT Collaborative Data Infrastructure, "B2STAGE - EUDAT," EUDAT Ltd, 2020. [Online]. Available: <https://www.eudat.eu/services/b2stage>. [Accessed 13 Apr. 2020].
- [7] M. D. Wilkinson et al., "The FAIR Guiding Principles for scientific data management and stewardship," *Scientific Data*, vol. 3, p. 160018, 2016.
- [8] N. Sakimura, J. Bradley, M. B. Jones, B. de Medeiros and C. Mortimore, "OpenID Connect Core 1.0 incorporating errata set 1," The OpenID Foundation, 8 Nov 2014. [Online]. Available: https://openid.net/specs/openid-connect-core-1_0.html. [Accessed 27 04 2020].
- [9] JBoss (Red Hat Inc.), Keycloak Community, "Keycloak," 2020. [Online]. Available: <https://www.keycloak.org/>. [Accessed 10 Apr. 2020].
- [10] M. Jones, B. Campbell and C. Mortimore, "RFC 7523 - JSON Web Token (JWT) Profile for OAuth 2.0 Client Authentication and Authorization Grants," Internet Engineering Task Force (IETF), May 2015. [Online]. Available: <https://tools.ietf.org/html/rfc7523>. [Accessed 1 Apr. 2020].
- [11] OASIS TC, "Topology and Orchestration Specification for Cloud Applications Version 1.0," OASIS Standard, 25 Nov. 2013. [Online]. Available: <http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.html>. [Accessed 27 Apr. 2020].
- [12] Atos BDS R&D, "Yorc 4.0.2," 2020. [Online]. Available: <http://yorc.readthedocs.io/>. [Accessed 29 Jul. 2020].
- [13] FastConnect, Bull/Atos, "Alien 4 Cloud," 2020. [Online]. Available: <http://alien4cloud.github.io/>. [Accessed 27 Apr. 2020].
- [14] V. Svatoň, J. Martinovič, J. Křenek, T. Esch and P. Tomančák, "HPC-as-a-Service via HEAppE Platform," in *Proceedings of the 13th International Conference on Complex, Intelligent, and Software Intensive Systems (CISIS-2019). Advances in Intelligent Systems and Computing*, 993, Sydney, Australia, 2019.
- [15] Skamarock, W. C. et al, "Technical Report NCAR/TN-556+STR, NCAR.," 2019.
- [16] Openstack Foundation, "Openstack," 2020. [Online]. Available: <https://www.openstack.org/>. [Accessed 01 Aug. 2020].
- [17] M. D'Andrea, P. Fiorucci, F. Gaetani and D. Negro, "RISICO: A decision support system (DSS) for dynamic wildfire risk evaluation in Italy," in *EGU General Assembly*, p.11102, Vienna, Austria, 2010.
- [18] Celery, "Celery: Distributed Task Queue," 2020. [Online]. Available: <http://www.celeryproject.org/>. [Accessed 1 Apr. 2020].
- [19] Django Software Foundation, "Django," 2020. [Online]. Available: <https://www.djangoproject.com/>. [Accessed 1 Apr. 2020].
- [20] Pivotal Software, "Messaging that just works — RabbitMQ," 2020. [Online]. Available: <https://www.rabbitmq.com/>. [Accessed 1 Apr. 2020].
- [21] Italian Civil Protection Department, & CIMA Research Foundation, „The Dewetra Platform: a multi-perspective architecture for risk management during emergencies.,“ in *First International Conference, ISCRAM-med 2014. Proceedings 1 (pp. 165-177)*, Toulouse, France, 2014.
- [22] PRACE, "Partnership For Advanced Computing in Europe," PRACE aisbl, Insight Media Group, 2020. [Online]. Available: <https://prace-ri.eu/>. [Accessed 01 Aug. 2020].