Contribution ID: **9**   Type: **Presentation (15' + 5' for questions)**

InterTwin's Digital Twin Engine Components OSCAR and DCNiOS

Tuesday, 29 October 2024 14:20 (20 minutes)

InterTwin co-designs and implements a prototype of an interdisciplinary Digital Twin Engine (DTE) - an open source platform based on open standards that offers the capability to integrate with application-specific Digital Twins (DTs).

While there are many components that are part of the DTE, this contribution focuses on OSCAR and DCNiOS and how they are being used in InterTwin to support the creation of DTs.

First, OSCAR is an open-source serverless framework that provides event-driven computing on scalable Kubernetes clusters. It supports the ability to run data processing containers in response to file uploads to an object store, as is the case of MinIO, a high-performance object storage that is deployed as part of an OSCAR cluster. InterTwin's data management layer involves the usage of dCache, a system for storing and retrieving huge amounts of data, distributed among a large number of heterogeneous server nodes, under a single virtual filesystem tree with a variety of standard access methods. To this aim, we integrated the ability to react to file upload events into dCache to trigger data processing inside the OSCAR cluster. For this, we created DCNiOS, a Data Connector between Apache NiFi and OSCAR, which facilitates the integration between the systems. The usage of NiFi allows us to buffer the data-processing requests coming from dCache to cope with the potentially different rates between data producer and data consumer. We integrated DCNiOS support for Apache Kafka as well to support Pub/Sub mechanisms for triggering the data processing.

To take advantage of hardware accelerators available in HPC facilities, we performed the integration between OSCAR and the INFN's interLink development, which provides a gateway to offload the execution of Kubernetes jobs into an HPC cluster. This allows to create an event-driven elastic computing platform which can offload the execution into HPC facilities, as was done for the HPC Vega supercomputer at IZUM, Slovenia.

In addition, the integration of itwinai, interTwin's AI platform for advanced AI/ML workflows in digital twin applications, with OSCAR, paved the way for an integrative approach to support general purpose event-driven computing with automated workload offloading, thus bridging Cloud and HPC facilities.

Finally, the ability to run generic services exposed in OSCAR clusters allows us to run Jupyter Notebooks that expose MinIO's storage system in the Notebook sandbox, thus facilitating data ingestion, processing, and visualization within the same platform.

This contribution will summarise the experiences in this area and lessons learnt during this process, as well as its applications to existing use cases in InterTwin.

Primary authors: PARCERO, Estíbaliz (Universitat Politècnica de València); LANGARITA, Sergio (Universitat Politècnica de València); ALARCÓN, Caterina (Universitat Politècnica de València); MOLTÓ, Germán (Universitat Politècnica de València)

Presenter: PARCERO, Estíbaliz (Universitat Politècnica de València)