## EVOLVE Webinars:

Bringing together HPC, Big Data and Cloud

Access Here

Evolve Newsletter Issue 5 May 2021

0

## EVOLVE is now entering in its final stretch, these last years we've addressed challenges, founds solutions, and also analyzed failures or missteps.

As in any serious attempts to innovate the journey has not been quiet and we've taken several technical risks, betting on emerging technologies, or at the opposite betting that some venerable old tech will not get disrupted. In this respect, EVOLVE has been quite successful containers are ubiquitous, Prometheus is the trendy solution for performance monitoring and Argo and Kubernetes are the de facto standard for workflow solutions. Overall, many results have been collected, alongside with some good memories too.

We are now proposing, out the successes and challenges encountered in EVOLVE a series of Webinar, not focused on the project itself but on some specific technologies we've been extensively used in the framework of EVOLVE.

This will be the opportunity to discuss during short session, always based on with hands-on examples, the articulation of Slurm with Kubernetes, workflow implementation, AI applied to agriculture optimization, Mobility with accelerated big data and many other topics.



Learn more about Evolve's implementation and POC's through our demos:

Running MPI NAS benchmarks on the EVOLVE platform

Watch Video

Running MPI NAS benchmarks on the EVOLVE platform

EVOLVE uses Kubernetes-based software stack with novel features. Applications are written in notebooks, using Apache Zeppelin. This short demo showcases the "MPI NAS Benchmarks" example notebook, which uses an MPI "virtual cluster" (a private collection of containers) to compile and execute one of the well-known NAS benchmarks. This example is meant to serve as a starting point for porting MPI applications to the EVOLVE platform. Moreover, it shows how to use the services API to start a new dashboard-based service from a notebook paragraph, and how to use automatically attached data namespaces (like "/remote") for interfacing with applications executing in containers.

– (Forth

Change detection pipeline on the EVOLVE platform

Watch Video



In this video, you will learn how to run Thales Alenia Space change detection tool from the EVOLVE frontend and understand what is exactly happening in the backend, in other words, how the different modules have been implemented, how pods are generated and how modules communicate between each other with Kafka.

ThalesAlenia

Vehicle assistance systems' status detection from video stream of a vehicle dashboard

Watch Video

Vehicle assistance systems' status detection from video stream of a vehicle dashboard

To better assess user interaction and acceptance of automated driving, knowledge of the status and status changes of driver assistance systems is essential. The demo video shows the execution of a workflow for status detection of vehicle assistance systems on the EVOLVE platform. This involves the use of computer vision methods for automatic state detection of two different driver assistance systems, lane assist and adaptive cruise control, using vehicle dashboard videos.

🕤 virtual 🛟 vehicle

End-to-end deployment capabilities of ML inference serving in EVOLVE through CI/CD pipeline

Watch Video

End-to-end deployment capabilities of ML inference serving in EVOLVE through CI/CD pipeline

This video shows the highly automation capabilities of EVOLVE's software stack in developing integrating and continuously deploying data and ML serving workloads on the EVOLVE platform.



Accelerated Network Anomaly Identification in EVOLVE through GPUenabled Spark deployments

Watch Video

Accelerated Network Anomaly Identification in EVOLVE through GPU-enabled Spark deployments

This video shows the GPU acceleration and custom monitoring capabilities of EVOLVE platform integrated with Apache SPARK data analytics engine. Through NUMBA Python JIT runtime, we integrate Apache Spark distributing management with NVIDIA GPUs and assign to GPU-featured nodes to accelerate the overall Spark job. The accelerated SPARK services are tightly integrated with the Zeppelin notebook and Graphana monitoring service environment to enable ease GPU acceleration to EVOLVE platform users.

1

Accelerating CNN Inference in EVOLVE platform with Intel Stratix10 FPGAs

Watch Video

Accelerating CNN Inference in EVOLVE platform with Intel Stratix10 FPGAs

This video shows the utilisation of FPGA acceleration capabilities for emerging CNN inference workloads deployed through

TensorFlow that forms a core functionality in EVOLVE dataanalytic use-cases. Although the FPGA acceleration forms a very high performant and power efficient solution it is not natively supported in the modern ML/AI frameworks. We show that within EVOLVE, FPGA acceleration it has been enabled and exposed through specific Python APIs to Zeppelin notebooks, allowing user and application developers to transparently benefit from the use of FPGA processing.



Latest developments on the Evolve's Platform

**Read article** 



During the previous periods, most of our developments have been dedicated to the convergence of the respective High-Performance Computing (HPC), Big Data (BD) and Machine Learning (ML) domains with high benefits expected from this cross-fertilization but raising at the same time multiple technical challenges. Indeed, HPC has been traditionally used for modeling and simulation while BD is dedicated to the ingestion and analysis of data from diverse sources. Both communities are evolving in response to changing user needs and technological landscapes and further with a convergence towards using machine learning (ML) not only for data analytics but also for modeling and simulation. These joint evolutions in the software field and at the hardware level put considerable pressure on system design. Evolve platform is the perfect vehicle to design and assess the converged architectures of tomorrow.

Author - Jean-Thomas <u>(DDN)</u> Huy-Nam Nguyen <u>(ATOS/BUL)</u>

## Be part of EVOLVE Ecosystem

The EVOLVE ecosystem aims to establish and nurture a network that will empower open innovation based on the platform that the evolve project is building. The EVOLVE ecosystem brings together all the policy makers, research organizations, engineers, developers and companies acting as end-users that are actively working on the fields of high-performance computing, big data and cloud.

The EVOLVE ecosystem will also facilitate innovative enterprises (large industries, SMEs, creative start-ups and university students – bachelor, master, PhD level) to develop and test their novel solutions.

Check the advantages of being part of the network and apply for membership



## Follow the project WWW.evolve-h2020.eu



