



OVERVIEW

Too often, data scientists and AI-engineers in an organization are some of the most precious resources. Naturally, it behooves the enterprise to do what it takes to improve their productivity and keep them happy. We thought not enough was being done in this regard, and we wanted to do something. In a similar vein, millions spent on grass-roots AI research in an academic setting is seeing limited success in terms of commercialization. At the same time, the enterprise data-scientists are eager to get their hands-on curated algorithms. So we started thinking about a fruit cart where one can go to sample different types of fruits, without going to different fruit farms. AI-Workstation is a realization of such a fruit cart.

Zeblok AI-Workstation builds upon the ubiquitous Jupyter Notebook technology and virtualizes it. It integrates a powerful and multi-class orchestration and scheduling layer to support a variety of workloads. A data-scientist, data-engineer, or an AI-engineer can now efficiently develop new AI models and take them to production with consummate ease. The sophisticated and yet simple orchestration layer supports both straightforward and non-computationally intensive AI-model-development, as well as the computationally-intensive workload. From an AI-workstation application, a user can utilize petaflops of computing power, with one simple and intuitive interface.

AI-Workstation provides enormous freedom to run enterprise workloads in many data centers based on cost profile and tailored to application requirements. This flexibility to drive their workloads is capital-efficient. A user can select compute resources from the University Data Centers, public clouds like AWS, or our business partner such as Atlantic Crypto Data Center.

Data scientists don't have to worry about optimizing frameworks on GPU platforms. CUDA optimization for popular AI frameworks is available out of the box and ready to go with a click of a button. Popular data science language bindings are also available such as R, Scala, and Python. We used a community of data scientists to beta-test and harden these containers so enterprise developers can spend more of their time developing innovative models. AI-Workstation also makes it easy for developers to build and share new containers with fellow team members improving organizational productivity. Dependency management, which is involved in AI-development, is greatly simplified and incorporated with examples for a developer to implement high-quality software. AI-Workstation also makes it a breeze to promote models developed into API through a few magic commands in the note.

Go ahead and confidently access petascale capacity through a few clicks from your favorite laptop.



FEATURES

ENORMOUS GPU CAPACITY

Start with 24 PFLOPS tensors (Machine Learning) through just a browser & self-registration, expanding to 200+ Petaflops in 2020.

FAMILIAR DATA SCIENCE TOOLS

Access ubiquitous data science tools like Jupyter Notebooks, Docker, Python, R, Scala make it a zero-learning-curve workstation for data scientists.

HIGHLY SECURE

Confidently use ActiveDirectory based Identity and Access Management in the environment that integrates web and OS level access controls into one platform.

MULTI-USER

Drive all your data scientists and engineers run workloads into one platform.

HIGH PERFORMANCE

Optionally run highspeed low latency interconnection through infiniband of up to 200 Gb/s

SUPERCOMPUTING

Create and distribute jobs across a network of GPUs, allowing cost-effective and speedy development of hyper-converged-AI applications.

API ENABLEMENT

Efficiently share trained engines via RESTful API to bring predictive insights into the enterprise.

NOVEL AI ALGORITHMS

Access to a curated list of AI-Algorithms, ready to be integrated into enterprise business processes.



BENEFITS:

- Powerful ML-AI notebook transforming your laptop into a Supercomputer.
- Standard Enterprise security integration enabled.
- Reduce time-to-market enormously when developing AI applications.
- Pay as you go lowers the development cost.

Zeblok Supercomputing AI-Workstation

