


PADS GPU server –Machine Learning and Data Ingestion for huge datasets

The Process and Data Science group provides access to a GPU server that can be used (after request and approval) for intensive computations related to the world of Machine Learning/Deep Learning/Computer Vision and to general, highly parallel, Process Mining computations. With the provision of a good amount of RAM, a connection with our Hadoop cluster for data storage, and threadripping CPUs/GPUs, no parallel computation is a fear.

Use Cases:


- Analysis of time series, both from process mining and generic logs, to forecast the future behavior
- Decision support from the historical data
- Computer vision (face/object detection; optical character recognition)
- Ingestion of CSV/Parquet log files for common data mining / process mining operations (calculation of the DFGs, filtering)

Advantages:



Query and store data, at a blazing speed

With more than 50 terabytes of disk space usable at a speed that is 10 times higher than the one of an high-performing notebook, every database/query can be easily managed in our cluster.




Smarter data preprocessing

The Process and Data Science Group is one of the leading data/process mining groups in the world. An important factor for the success of a data science project is the correct preprocessing of the data.



Magnify your process with the prediction lens

Our process mining group is researching and developing prediction/recommendation techniques. Predict the outcome/remaining time of your process executions, training your models on our advanced platform.



GPU Data Ingestion – The dawn of a new era

In the previous years, GPUs were relegated to Machine Learning tasks. Follow the rise of new paradigms to load huge amount of data in-memory and process them at a blazing speed.