On-Demand Distributed Compute

Develop Complex Models Faster

Overview

Distributed software frameworks and compute clusters help data science teams solve the most complex machine learning problems — those that require large amounts of data and enormous processing power.

However, manually configuring clusters takes specialized DevOps skills and can be very time-consuming. That's why Domino provides self-service access to the three most popular distributed compute frameworks – Spark, Ray, and Dask. Teams can select the best framework for the job at hand, and quickly set up clusters, so they can test more ideas and develop better models faster.

Domino automatically scales compute clusters based on the workload to simplify provisioning, optimize utilization, and manage computing costs so you can maximize the productivity of your teams and of the return on your computing investments.

Benefits

- **Develop better models faster:** Put more powerful infrastructure at data scientists' fingertips to they can test many more ideas and solve more complex problems.
- **Easy to use for data scientists:** Provide on-demand access to distributed compute, and simplify package and dependency management so your data scientists can focus on data science.
- **Achieve significant cost savings:** Optimize resource utilization and the overall Domino cluster sizes that needsto be active an any one time with auto-scaling, so more people can share your infrastructure.
- **Easy to manage and support for IT:** Reduce the amount of IT support needed with centralized configuration of cluster settings, and self-service cluster provisioning that adheres to networking and security rules.

On-Demand Distributed Compute Features

Choice of Distributed Frameworks

- Spin up a cluster with just a few clicks — No DevOps experience required! Domino automatically connects the cluster to your workspace (e.g., Jupyter) or batch script.
- Provide the flexibility to leverage three of the most popular frameworks for distributed computing, so data scientists can use the best approach for a given use case.
- Take advantage of Ray and Dask, and their integrations with data science libraries, to easily leverage clusters for your Python-based data science code with minimal adaptation.

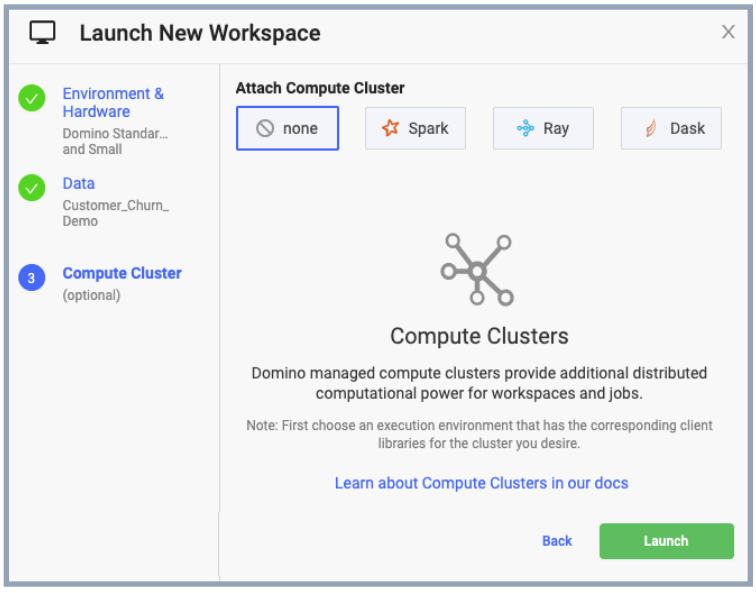

Self-Service Provisioning

- No need to wait on others to provision hardware — simply select the appropriate hardware tiers.
- Use multiple CPUs, GPUs, or large memory pools and speed up training.
- Automatically de-provision clusters when work is completed, saving the organization costs and freeing up resources for others.

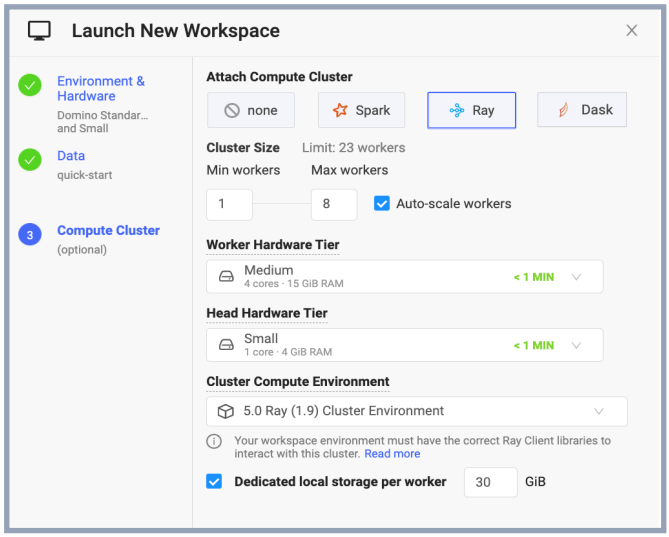

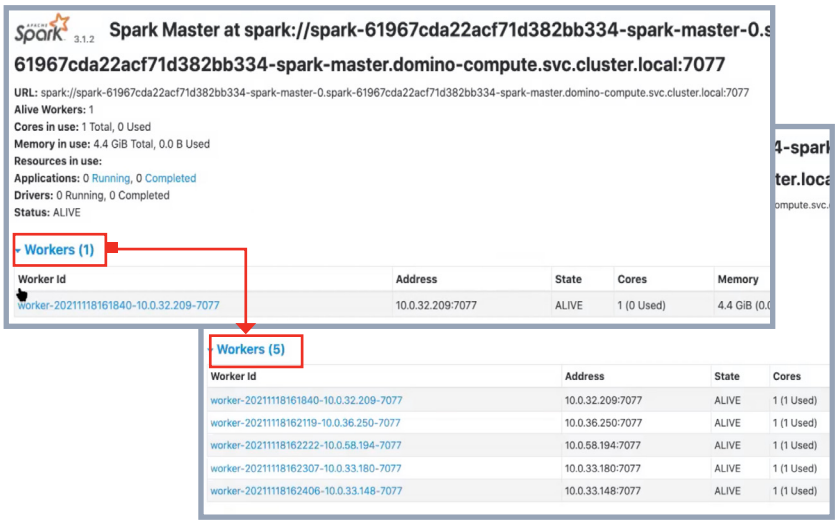

Autoscaled Compute

- The autoscaling option allows clusters to automatically adpat to workloads.
- Eliminate the need to estimate and commit resources ahead of time.
- Optimize your consumption of compute so other members of your team can share resources.

Simplified IT Management

- Single footprint to manage — no need for multiple clusters or ad hoc requests.
- Painless package management and governance with versioned cluster compute environments.
- Specify specific hardware tiers for distributed compute frameworks to optimize hardware utilization.

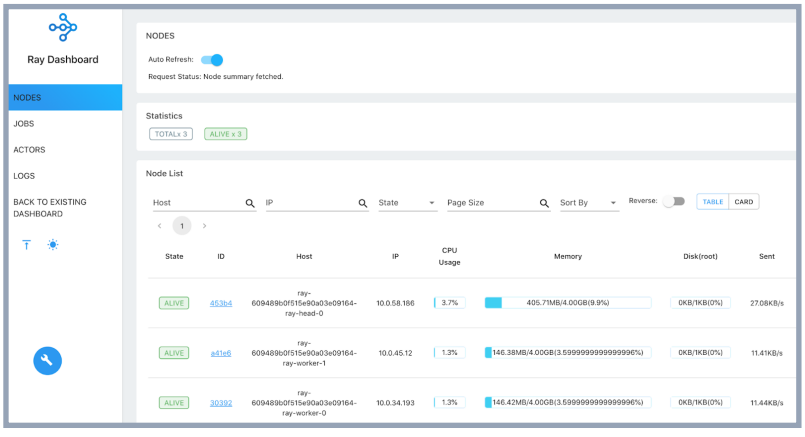

- With On-Demand Distributed Compute, the Domino Enterprise MLOps platform puts self-service infrastructure at data scientists' fingertips with full reproducibility and infrastructure governance.