ANACONDA.

Dask Infrastructure
Kickstart

Easily run large distributed compute jobs on your existing infrastructure

The Dask Infrastructure Kickstart will enable you to run large compute jobs on your existing Kubernetes, Hadoop, or HPC resources, on-prem or in the cloud, without having to rewrite your Python-based workflows. Our Dask experts help you install and configure Dask Gateway to access your CPU and GPU compute resources and to integrate with Jupyter Notebooks. We also help you develop a custom strategy for managing environments to ensure Dask-related local and remote libraries match.

In a short 100-hour engagement, our experts help you develop a comprehensive plan to get the most from the unique infrastructure you already have in place - without sacrificing efficiency or administrative control.

Requirements

To qualify for this Kickstart project, you must be able to provide the following:

- An overview of the types of workloads you are looking to parallelize, to make sure that we configure the system appropriately

- Descriptions of the compute resources you want to make available to Dask, along with administrative contacts who can configure them to accept Dask tasks at our direction

- Access to sample data sources for the Dask workers to process

- License for Anaconda Enterprise Edition, Enterprise 4, or Anaconda Team Edition

Get started using Dask with three deliverables:

Dask-gateway installation, including a deployment script and keep-alive example, with a backend configured to your specifications

Integration with Jupyter notebooks (in Anaconda Enterprise or via JupyterHub) to allow launching and monitoring multi-worker jobs interactively from notebooks

Customized strategy for managing your environments to ensure Dask-related local and remote libraries match

Get started today

1 Confirm that Dask Infrastructure Kickstart is a good fit for your organization.

2 Contact an Anaconda Kickstart specialist to discuss your project.

3 Start using Dask for distributed computing on your hardware!