

# BRIGHT CLUSTER MANAGER FOR DATA SCIENCE

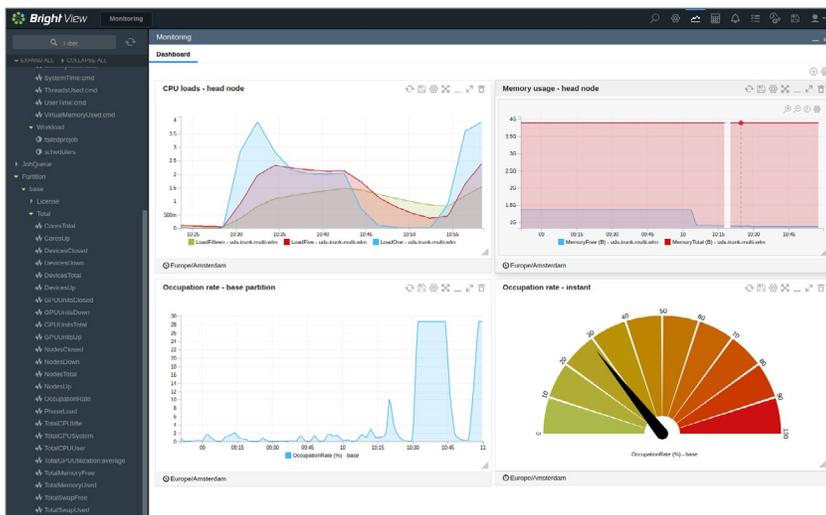
## Deep Learning Infrastructure Management Made Easy



For organizations that need to easily deploy and manage data science clusters

Bright Cluster Manager for Data Science is an add-on to Bright Cluster Manager—providing the computing infrastructure and pre-packaged tools you need to accelerate your data science projects and make them production-ready.

Bright brings everything together in a simple, easy to deploy and manage solution that will get you up and running quickly and reliably and with confidence.



Monitor every aspect of your data science cluster at a glance

# EASY TO

DEPLOY  
OPERATE  
MONITOR  
MANAGE  
SCALE

## Features

### Modern Deep Learning Environment

Provides everything you need to quickly spin up a complete machine learning environment for data scientists that is cluster-ready and easy to manage

### Choice of Machine Learning Libraries and Frameworks

Includes the most popular machine learning libraries and frameworks such as NVIDIA cuDNN, CUB, CUDA, TensorRT, Dynet, Fastai, JupyterHub, NCCL2, MXNet, pyTorch, Chainer, CNTK, Horovod, Keras, Opencv3, Protobuf3, Torch, Tensorflow, Theano, and XGBoost (see: <https://support.brightcomputing.com/packages-dashboard/>).

### Scale-out JupyterHub

JupyterHub makes data science easy to use. Bright provides custom JupyterHub spawners that allow notebooks to be scheduled through the HPC scheduler or Kubernetes.

### Support for NVIDIA GPU Cloud (NGC)

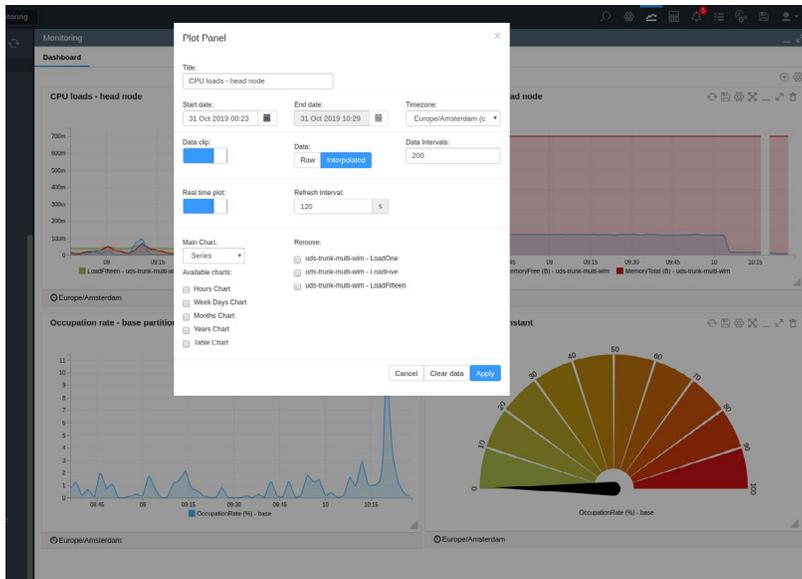
Makes it easy to use NVIDIA NGC deep learning containers in so many ways—run them in containers, through a batch scheduler, on physical nodes, or in a cloud

Enhances job metrics, accounting and reporting functionality making it possible to determine GPU usage on a per-job basis over time

### Simple Deployment Process

Installs all the software and creates the necessary configuration files by simply answering a few questions about your cluster and Bright takes care of the rest

Builds your new cluster right the first time and every time—on bare metal or in the cloud.



Simply drag the component you want to monitor into the display area and Bright Cluster Manager instantly creates a graph of the data for you.

Bright's user interface makes it easy to monitor and manage your data science clusters.

Customizable dashboards help administrators ensure all nodes in a cluster are optimized for the workloads they run.

Bright for Data Science makes it faster and easier for you to gain actionable insights from rich, complex data by combining popular machine learning and deep learning software packages with fully automated and scalable infrastructure for data scientists.

## Build a cost-efficient Spark cluster

Looking for a lightning-fast big data cluster solution? Get the maximum performance for your big data analytics using Bright Cluster Manager to deploy, monitor, and manage your own Spark cluster. You will receive all the benefits of Spark's fast, in-memory processing, along with Bright's famous ease-of-use capabilities.

## Your cloud, your way

Bright Cluster Manager can provision and manage clusters running on virtual servers inside a public cloud as if they were local machines. Use this feature to build an entire cluster in Amazon Web Services (AWS) or Microsoft® Azure from scratch, or extend a physical cluster into the cloud when you need extra capacity. And as AWS and Azure are now integrated with Bright, usage of Bright in the cloud is billed on an hourly basis and included on the monthly statement from your cloud provider. An added benefit of Bright Cluster Manager is Data-set Labeling, which enables you to create pipelines of jobs where Job A produces output that will serve as input for Job B, and all data is sent to/from the cloud only once.

## Key Benefits

### Simple deployment of deep learning infrastructure

– With Bright, you can deploy deep learning environments in minutes. Bright simplifies management of multiple-architecture and multiple-OS clusters by combining different types of servers, architectures, and OSes in the same cluster.

**Developer-friendly environment** – Bright for Data Science provides and installs environment modules that make it easy to dynamically modify the user environment. The software comes with a complete set of tools and libraries for HPC, Deep Learning, and Apache Spark so you will be ready to develop, debug, and deploy your code immediately.

**Scalability** – Bright makes HPC technology available to machine learning frameworks, allowing deep learning workloads to be spread across multiple machines. Bright's new Auto-Scale feature automatically scales HPC queues and Kubernetes applications from pools of both local nodes and cloud nodes to create a highly effective hybrid cloud environment.

**Extended capabilities** – Need more capacity? Extend your GPU-enabled instances into the cloud using Bright's cloud bursting capability. Bright also makes it easy to use NVIDIA machine learning containers (NGC) in your own private OpenStack cloud. In addition, Bright OpenStack includes the Stein release of OpenStack and the Nautilus release of Ceph to provide multi-attach for RDB volumes, optional deferred asynchronous volume deletion, bandwidth-aware VM scheduling, and more.