

---

# Data Science Stack

Canonical Group Ltd

Apr 07, 2025



# CONTENTS

<b>1</b>	<b>In this documentation</b>	<b>3</b>
<b>2</b>	<b>Project and community</b>	<b>5</b>
2.1	Tutorial . . . . .	5
2.2	How-to guides . . . . .	7
2.3	Explanation . . . . .	18



Data Science Stack (DSS) is a ready-to-run environment for Machine Learning (ML) and data science. It's built on open-source tooling, including Canonical K8s, JupyterLab, and MLflow, and is usable on any Ubuntu/Snap-enabled workstation.

DSS provides a Command Line Interface (CLI) for managing containerised ML environment images such as PyTorch or TensorFlow, on top of Canonical K8s.

Typically, creating ML environments on a workstation involves complex and hard-to-reverse configurations. DSS solves this problem by providing accessible, production-ready, isolated, and reproducible ML environments that fully utilise a workstation's GPUs.

Both ML beginners and engineers who need to build complex development and runtime environments will see set-up time reduced to a minimum, allowing them to get started with meaningful work within minutes.

---



## IN THIS DOCUMENTATION

### *Tutorial*

Get started - a hands-on introduction to DSS for newcomers

### *How-to guides*

Step-by-step guides covering key operations and common tasks with DSS

### *Explanation*

Discussion and clarification of key topics

---





## PROJECT AND COMMUNITY

Data Science Stack is an open-source project that values its community. We warmly welcome contributions, suggestions, fixes, and constructive feedback from everyone.

- Read our [Code of conduct](#).
- [Contribute](#) and [report bugs](#).
- Join the [Discourse forum](#).
- Talk to us on [Matrix](#).

### 2.1 Tutorial

This section of our documentation contains a step-by-step tutorial to get you started with Data Science Stack (DSS). You will learn how it works, how to customise your setup and explore what you can do with it:

#### 2.1.1 Get started with DSS

This guide describes how you can get started with Data Science Stack (DSS). From setting up Canonical K8s in your host environment, all the way to running your first notebook.

Data Science Stack is a ready-made environment that makes it seamless to run GPU-enabled containerised Machine Learning (ML) environments. It provides easy access to a solution for developing and optimising ML models, utilising your machine's GPUs and allowing users to utilise different ML environment images based on their needs.

#### Requirements

- Ubuntu 24.04 LTS.
- [Snap](#) installed.
- 50GB of disk space is recommended.

#### Set up Canonical K8s

DSS relies on a container orchestration system, capable of exposing the host GPUs to the workloads. [Canonical K8s](#) is used as the orchestration system. All the workloads and state managed by DSS are running on top of Canonical K8s.

You can install Canonical K8s using `snap` as follows:

```
sudo snap install k8s --classic --channel=1.32-classic/stable
sudo k8s bootstrap
sudo k8s enable local-storage
```

### Install the DSS CLI

DSS is a Command Line Interface (CLI)-based environment. Now, install the DSS CLI using the following command:

```
sudo snap install data-science-stack
```

### Initialise DSS

Next, you need to initialise DSS on top of Canonical K8s and prepare MLflow:

```
dss initialize --kubeconfig="$(sudo k8s config)"
```

#### Note

The initialisation might take a few minutes to complete. While the process is in progress, you will see the following message: `Waiting for deployment mlflow in namespace dss to be ready...`

Once initialised, you should see an output like this:

```
Deployment mlflow in namespace dss is ready
```

### Launch a Notebook

At this point, DSS is set up and you are ready to start managing containerised notebook environments. You can use the DSS CLI to launch and access them using JupyterLab.

Start your first Jupyter Notebook with the following command:

```
dss create my-tensorflow-notebook --image=kubeflownotebookswg/jupyter-tensorflow-cuda:v1.8.0
```

Once the command succeeds, it returns a URL that can be used to connect to the JupyterLab User Interface (UI) of that notebook. For example, you should expect an output like this:

```
Executing create command
Waiting for deployment my-tensorflow-notebook in namespace dss to be ready...
Deployment my-tensorflow-notebook in namespace dss is ready
Success: Notebook my-tensorflow-notebook created successfully.
Access the notebook at http://10.152.183.42:80.
```

Once you know the URL, open a web browser and enter the URL into the address bar. This will direct you to the notebook UI where you can start working with your notebook.

### Next Steps

- To learn more about how to interact with DSS, see [Manage DSS](#).
- To connect to MLflow, see [Manage MLflow](#).
- To leverage your GPUs, see [Enable GPUs](#).

## 2.2 How-to guides

The following guides cover key processes in the Data Science Stack (DSS) environment. Specifically, they demonstrate how to perform common tasks for managing DSS components.

### 2.2.1 DSS basics

Learn basic operations to interact with DSS, including installation and the Command Line Interface (CLI) commands:

#### Manage DSS

This guide describes how to manage Data Science Stack (DSS).

DSS is a Command Line Interface (CLI)-based environment and distributed as a [snap](#).

#### Install

##### **Note**

To install DSS, ensure you have previously installed [Snap](#) and [Canonical K8s](#).

You can install DSS using `snap` as follows:

```
sudo snap install data-science-stack
```

Then, you can run the DSS CLI with:

```
dss
```

#### Initialise

You can initialise DSS through `dss initialize`. This command:

- Stores credentials for the Canonical K8s cluster.
- Allocates storage for your DSS Jupyter Notebooks.
- Deploys an [MLflow](#) model registry.

```
dss initialize --kubeconfig "$(sudo k8s config)"
```

The `--kubeconfig` option is used to provide your Canonical K8s cluster's kubeconfig.

##### **Note**

Note the use of quotes for the `--kubeconfig` option. Without them, the content may be interpreted by your shell.

You should expect an output like this:

```
Executing initialize command
Storing provided kubeconfig to /home/user/.dss/config
Waiting for deployment mlflow in namespace dss to be ready...
Deployment mlflow in namespace dss is ready
```

(continues on next page)

(continued from previous page)

DSS initialized. To create your first notebook run the command:

```
dss create
```

Examples:

```
dss create my-notebook --image=pytorch
```

```
dss create my-notebook --image=kubeflownotebookswg/jupyter-scipy:v1.8.0
```

### Remove

You can remove DSS from your Canonical K8s cluster through `dss purge`. This command purges all the DSS components, including:

- All Jupyter Notebooks.
- The MLflow server.
- Any data stored within the DSS environment.

#### Note

This action removes the components of the DSS environment, but it does not remove the DSS CLI or your Canonical K8s cluster. To remove those, [delete their snaps](#).

```
dss purge
```

#### Caution

This action is irreversible. All data stored within the DSS environment will be lost.

You should expect an output like this:

```
Waiting for namespace dss to be deleted...
Success: All DSS components and notebooks purged successfully from the Kubernetes_
↳cluster.
```

### Get status

You can check the DSS status through `dss status`. This command provides a quick way to check the status of your DSS environment, including the MLflow status and whether a GPU is detected in your environment.

```
dss status
```

If you already have a DSS environment running and no GPU available, the expected output is:

```
MLflow deployment: Ready
MLflow URL: http://10.152.183.68:5000
GPU acceleration: Disabled
```

## List commands

You can get the list of available commands for DSS through the `dss` command with the `--help` option:

```
dss --help
```

You should expect an output like this:

```
Usage: dss [OPTIONS] COMMAND [ARGS]...

Command line interface for managing the DSS application.

Options:
--help  Show this message and exit.

Commands:
create  Create a Jupyter notebook in DSS and connect it to MLflow.
initialize  Initialize DSS on the given Kubernetes cluster.
list    Lists all created notebooks in the DSS environment.
logs    Prints the logs for the specified notebook or DSS component.
purge   Removes all notebooks and DSS components.
remove  Remove a Jupyter Notebook in DSS with the name NAME.
start   Starts a stopped notebook in the DSS environment.
status  Checks the status of key components within the DSS...
stop    Stops a running notebook in the DSS environment.
```

### Get details about a specific command:

To see the usage and options of a DSS command, run `dss <command>` with the `--help` option. For example:

```
dss logs --help
```

You should expect an output like this:

```
Usage: dss logs [OPTIONS] [NOTEBOOK_NAME]

Prints the logs for the specified notebook or DSS component.

Examples:
  dss logs my-notebook
  dss logs --mlflow
  dss logs --all

Options:
--kubeconfig TEXT  Path to a Kubernetes config file. Defaults to the value
                   of the KUBECONFIG environment variable, else to
                   './kubeconfig'.
--all              Print the logs for all notebooks and MLflow.
--mlflow           Print the logs for the MLflow deployment.
--help            Show this message and exit.
```

### See also

- To learn how to manage your Jupyter Notebooks, check *Manage Jupyter Notebooks*.
- If you are interested in managing MLflow within your DSS environment, see *Manage MLflow*.

## 2.2.2 Jupyter Notebooks within DSS

Learn how to manage [Jupyter Notebooks](#) within your DSS environment:

### Manage Jupyter Notebooks

This guide describes how to manage [Jupyter Notebooks](#) within your Data Science Stack (DSS) environment.

All actions can be performed using the DSS Command Line Interface (CLI).

### Create a notebook

You can create a Jupyter Notebook using the DSS CLI. This notebook includes different packages and toolkits depending on the image used to create it.

#### 1. Select an image:

Before creating a notebook, you need to select an image that includes the packages and toolkits you need. To see a list of recommended images and their aliases, do:

```
dss create --help
```

The output includes a list of recommended images and their aliases. For example, this guide uses the image *kubeflownotebookswg/jupyter-scipy:v1.8.0*

#### 2. Create the notebook:

Create a new notebook as follows:

```
dss create my-notebook --image kubeflownotebookswg/jupyter-scipy:v1.8.0
```

This command starts a notebook server with the selected image. You should expect an output like this:

```
Executing create command
Waiting for deployment test-notebook in namespace dss to be ready...
Deployment test-notebook in namespace dss is ready
Success: Notebook test-notebook created successfully.
Access the notebook at http://10.152.183.42:80.
```

### Create an NVIDIA GPU-enabled notebook

You can create an NVIDIA GPU-enabled Jupyter Notebook containing CUDA runtimes and Machine Learning (ML) frameworks, and access its JupyterLab server.

#### Note

To launch an NVIDIA GPU-enabled notebook, you must first *install* the NVIDIA Operator and *verify* DSS can detect the GPU. See *Enable NVIDIA GPUs* for more details.

To see the list of available CUDA images, run:

```
dss create --help | grep cuda
```

You should see an output similar to this:

```
- pytorch-cuda = kubeflownotebookswg/jupyter-pytorch-cuda-full:v1.8.0
- tensorflow-cuda = kubeflownotebookswg/jupyter-tensorflow-cuda-full:v1.8.0
```

Select one of them and create a notebook as follows:

```
dss create my-notebook --image=tensorflow-cuda
```

You can confirm your GPU is detected and usable by running the following within your notebook:

```
import tensorflow as tf

tf.config.list_physical_devices('GPU')
```

### Create an Intel GPU-enabled notebook

You can create an Intel GPU-enabled Jupyter Notebook with [Intel Extension for PyTorch \(IPEX\)](#) or [Intel Extension for TensorFlow \(ITEX\)](#).

#### **Note**

To launch an Intel GPU-enabled notebook, you must first *Enable Intel GPUs*.

To see the list of available Intel images, run:

```
dss create --help | grep intel
```

You should see an output similar to this:

```
- pytorch-intel= intel/intel-extension-for-pytorch:2.1.20-xpu-idp-jupyter
- tensorflow-intel = intel/intel-extension-for-tensorflow:2.15.0-xpu-idp-jupyter
```

Select one of them and create a notebook as follows:

```
dss create my-itex-notebook --image=tensorflow-intel
```

You can confirm your Intel GPU is detected and usable by running the following within your notebook:

```
import tensorflow as tf

tf.config.experimental.list_physical_devices()
```

### List created notebooks

You can check the current state of all notebooks within your DSS environment. To view the full list, run:

```
dss list
```

This command displays each notebook name along with its associated image, state, and URL if applicable.

### Remove a notebook

You can remove a Jupyter Notebook using the DSS CLI. It is a non-blocking process, meaning you can continue other work while the deletion completes.

1. **Remove the notebook:**

To delete the notebook, use the `dss remove` command followed by the name of the notebook, `my-notebook` in this example:

```
dss remove my-notebook
```

You should expect an output like this:

```
Removing the notebook my-notebook. Check `dss list` for the status of the notebook.
```

### Start a notebook

You can start a notebook using the DSS CLI. This enables you to resume your work without needing to configure a new notebook.

```
dss start my-notebook
```

### Stop a notebook

You can stop a notebook using the DSS CLI. Stopping a notebook frees up resources and ensures data safety when not actively working on it.

```
dss stop my-notebook
```

### Access a notebook

You can access a notebook User Interface (UI) using the DSS CLI. Accessing the UI enables you to interact directly with your notebook, run code, and visualise data. This is done through a web browser by navigating to the URL associated with your active notebook.

1. **Find the notebook URL:**

To find the URL of your notebook, first list all the notebooks:

```
dss list
```

Look for your notebook in the output, and specifically check the URL column.

2. **Access the Notebook UI:**

Once you know the URL, open a web browser and enter the URL into the address bar.

### Get notebook logs

You can retrieve logs for a Jupyter Notebook using the DSS CLI. Retrieving logs can help you troubleshoot issues, monitor notebook activities, or verify actions taken in the notebook.

```
dss logs my-notebook
```



## See also

- To learn how to manage your DSS environment, check *Manage DSS*.
- If you are interested in managing MLflow within your DSS environment, see *Manage MLflow*.

## 2.2.3 MLflow within DSS

Learn how to manage [MLflow](#) within your DSS environment:

### Manage MLflow

This guide describes how to manage [MLflow](#) within your Data Science Stack (DSS) environment.

MLflow is a platform for managing the end-to-end machine learning life cycle. It includes tracking experiments, packaging code into reproducible runs, and sharing and deploying models.

### Access MLflow

You can access the MLflow User Interface (UI) within your DSS environment through a web browser, by navigating to the URL associated with MLflow. This UI allows you to interact directly with your MLflow experiments and models.

#### 1. Get the MLflow URL:

To find the URL of MLflow, run:

```
dss status
```

Look for the MLflow URL in the output. For example:

```
MLflow deployment: Ready
MLflow URL: http://10.152.183.205:5000
```

#### **i** Note

To access the UI, your MLflow deployment should be *Ready*.

#### 2. Access the MLflow UI:

Once you know the URL, open a web browser and enter the URL into the address bar. This will direct you to the MLflow interface.

### Get logs

You can retrieve MLflow logs within your DSS environment. Retrieving logs is a critical task for maintaining and troubleshooting MLflow.

To get MLflow logs, use the `dss logs` command with the `--mlflow` option:

```
dss logs --mlflow
```

You should expect an output like this:

```
Logs for mlflow-6bbfc5db5-xlfvj:
[2024-04-30 07:57:54 +0000] [22] [INFO] Starting gunicorn 20.1.0
[2024-04-30 07:57:54 +0000] [22] [INFO] Listening at: http://0.0.0.0:5000 (22)
```

(continues on next page)

(continued from previous page)

```
[2024-04-30 07:57:54 +0000] [22] [INFO] Using worker: sync
[2024-04-30 07:57:54 +0000] [23] [INFO] Booting worker with pid: 23
[2024-04-30 07:57:54 +0000] [24] [INFO] Booting worker with pid: 24
[2024-04-30 07:57:54 +0000] [25] [INFO] Booting worker with pid: 25
[2024-04-30 07:57:54 +0000] [26] [INFO] Booting worker with pid: 26
```

### Get artefacts

MLflow artefacts, including [models](#), [experiments](#) and [runs](#) are stored within your DSS environment. You can access and download them from the [MLflow UI](#).

### See also

- To learn how to manage your DSS environment, check [Manage DSS](#).
- If you are interested in managing Jupyter Notebooks within your DSS environment, see [Manage Jupyter Notebooks](#).
- See [Charmed MLflow](#) for more details on MLflow.

## 2.2.4 Enable GPUs

Learn how to configure DSS to leverage your GPUs:

### Enable GPUs

The following guides cover configuration aspects to leverage your GPUs within the Data Science Stack (DSS) environment.

### Enable Intel GPUs

This guide describes how to configure Data Science Stack (DSS) to utilise the Intel GPUs on your machine.

You can do so by enabling the Intel device plugin on your [Canonical K8s](#) cluster.

### Prerequisites

- Ubuntu 24.04.
- DSS is *installed* and *initialised*.
- The `kubectl snap` package is installed.
- Your machine includes an Intel GPU.

### Verify the Intel GPU drivers

To confirm that your machine has the Intel GPU drivers set up, first install the `intel-gpu-tools` package:

```
sudo apt install intel-gpu-tools
```

Now list the Intel GPU devices on your machine as follows:

```
intel_gpu_top -L
```

If the drivers are correctly installed, you should see information about your GPU device such as the following:

```
card0                8086:56a0
pci:vendor=8086,device=56A0,card=0
└─renderD128
```

### Note

For Intel discrete GPUs on Ubuntu versions older than 24.04, you may need to perform additional steps such as installing a [HWE kernel](#).

## Enable the Intel GPU plugin

To ensure DSS can utilise Intel GPUs, you have to enable the Intel GPU plugin in your Canonical K8s cluster.

1. Use `kubectl kustomize` to build the plugin YAML configuration files:

```
VERSION=v0.30.0
kubectl kustomize https://github.com/intel/intel-device-plugins-for-kubernetes/
↳ deployments/nfd?ref=${VERSION} > node_feature_discovery.yaml
kubectl kustomize https://github.com/intel/intel-device-plugins-for-kubernetes/
↳ deployments/nfd/overlays/node-feature-rules?ref=${VERSION} > node_feature_rules.yaml
kubectl kustomize https://github.com/intel/intel-device-plugins-for-kubernetes/
↳ deployments/gpu_plugin/overlays/nfd_labeled_nodes?ref=${VERSION} > gpu_plugin.yaml
```

To allow multiple containers to utilise the same GPU, run:

```
sed -i 's/enable-monitoring/enable-monitoring\n          - --shared-dev-num=10/' gpu_plugin.
↳ yaml
```

2. Apply the built YAML files to your Canonical K8s cluster:

```
kubectl apply -f node_feature_discovery.yaml
kubectl apply -f node_feature_rules.yaml
kubectl apply -f gpu_plugin.yaml
```

The Canonical K8s cluster is now configured to recognise and utilise your Intel GPU.

### Note

After the YAML configuration files have been applied, they can be safely deleted.

## Verify the Intel GPU plugin

To verify the Intel GPU plugin is installed and the Canonical K8s cluster recognises your GPU, run:

```
kubectl get nodes --show-labels | grep intel
```

You should see an output with the cluster name such as the following:

```
kubectl get nodes --show-labels | grep intel
fluent-greenshank Ready <none> 18s v1.30.3 beta.kubernetes.io/arch=amd64,beta.
↳ kubernetes.io/os=linux,intel.feature.node.kubernetes.io/gpu=true
```

### Verify DSS detects the GPU

Verify DSS has detected the GPU by checking the DSS status. To do so, run the following command using the DSS CLI:

```
dss status
```

You should expect an output like this:

```
Output:
MLflow deployment: Ready
MLflow URL: http://10.152.183.68:5000
NVIDIA GPU acceleration: Disabled
Intel GPU acceleration: Enabled
```

### See also

- To learn how to manage your DSS environment, check [Manage DSS](#).
- If you are interested in managing Jupyter Notebooks within your DSS environment, see [Manage Jupyter Notebooks](#).

### Enable NVIDIA GPUs

This guide describes how to configure Data Science Stack (DSS) to utilise your NVIDIA GPUs within a Canonical K8s environment.

DSS supports GPU acceleration by leveraging the [NVIDIA GPU Operator](#). The operator ensures that the necessary components, including drivers and runtime, are set up correctly to enable GPU workloads.

### Prerequisites

- DSS is *installed* and *initialised*.
- Your machine includes an NVIDIA GPU.

### Install the NVIDIA GPU Operator

To enable GPU support, you must install the NVIDIA GPU Operator in your Kubernetes cluster. Follow [NVIDIA GPU Operator Installation Guide](#) for installation details.

### Verify the NVIDIA Operator is up

Once the NVIDIA GPU Operator is installed, verify that it has been successfully initialized before running workloads.

### Ensure DaemonSet is ready

Run the following command to verify that the DaemonSet for the NVIDIA Operator Validator is created:

```
while ! kubectl get ds -n gpu-operator-resources nvidia-operator-validator; do
  sleep 5
done
```

**Note**

It may take a few seconds for the DaemonSet to be created.

Once completed, you should see an output similar to this:

NAME	DESIRED	CURRENT	READY	UP-TO-DATE	AVAILABLE	NODE_
↪SELECTOR			AGE			
nvidia-operator-validator	1	1	0	1	0	nvidia.
↪com/gpu.deploy.operator-validator=true			17s			

**Ensure the Validator Pod succeeded**

Run the following command to check if the NVIDIA Operator validation is successful:

```
echo "Waiting for the NVIDIA Operator validations to complete..."

while ! kubectl logs -n gpu-operator-resources -l app=nvidia-operator-validator -c_
↪nvidia-operator-validator | grep "all validations are successful"; do
  sleep 5
done
```

**Note**

If the pod is still initializing, you may see an error like: Error from server (BadRequest): container "nvidia-operator-validator" in pod "nvidia-operator-validator-xxxx" is waiting to start: PodInitializing

Once completed, the output should include:

```
all validations are successful
```

**Use cases for different driver states**

The NVIDIA GPU Operator behaves differently depending on whether your system already has an NVIDIA driver installed. Below are the three primary scenarios:

**Device with no NVIDIA driver installed**

- The GPU Operator will **automatically install** the necessary NVIDIA driver.
- This installation process may take longer as it involves setting up drivers and runtime components.
- Once the process is complete, GPU should be detected successfully.

**Device with an up-to-date NVIDIA driver**

- The GPU Operator detects the existing driver and proceeds without reinstalling it.
- However, to avoid redundant installations, it is recommended to disable driver installation explicitly when deploying the operator.
- Follow the upstream documentation for the correct configuration to disable driver installation in this case.

### Device with an outdated NVIDIA driver

- If an older driver version is detected, the operator may attempt to install a newer version.
- This could lead to conflicts if the outdated driver does not match the required CUDA version.
- To prevent issues, update the driver manually or remove the outdated version before deploying the GPU Operator.

### Verify DSS detects the GPU

After installing and configuring the NVIDIA GPU Operator, verify that DSS detects the GPU by checking its status:

```
dss status
```

You should expect an output like this:

```
MLflow deployment: Ready
MLflow URL: http://10.152.183.74:5000
GPU acceleration: Enabled (NVIDIA-GeForce-RTX-3070-Ti)
```

#### Note

The GPU model displayed may differ based on your hardware.

### See also

- To learn how to manage your DSS environment, check [Manage DSS](#).
- If you are interested in managing Jupyter Notebooks within your DSS environment, see [Manage Jupyter Notebooks](#).

## 2.3 Explanation

The following guides cover key concepts and features of Data Science Stack (DSS).

### 2.3.1 Architecture overview

Understand the underlying architecture of DSS, its components and interactions:

#### DSS architecture

This guide provides an overview of the Data Science Stack (DSS) architecture, its main components, and their interactions.

DSS is a ready-to-run environment for Machine Learning (ML) and Data Science (DS). It's built on open-source tooling, including [Canonical K8s](#), [JupyterLab](#), and [MLflow](#).

DSS is distributed as a [snap](#) and usable on any Ubuntu workstation. This provides robust security management and user-friendly version control, enabling seamless updates and auto-rollback in case of failure.

Using DSS, you can perform the following tasks:

- Installing and managing the DSS Python library.
- Deploying and managing Jupyter Notebooks.
- Deploying and managing MLflow.

- Running GPU workloads.

### Architecture overview

The DSS architecture can be thought of as a stack of layers. These layers, from top to bottom, include:

- *Application.*
- *ML tools.*
- *Orchestration.*
- *Operating system (OS).*

The following diagram showcases it:

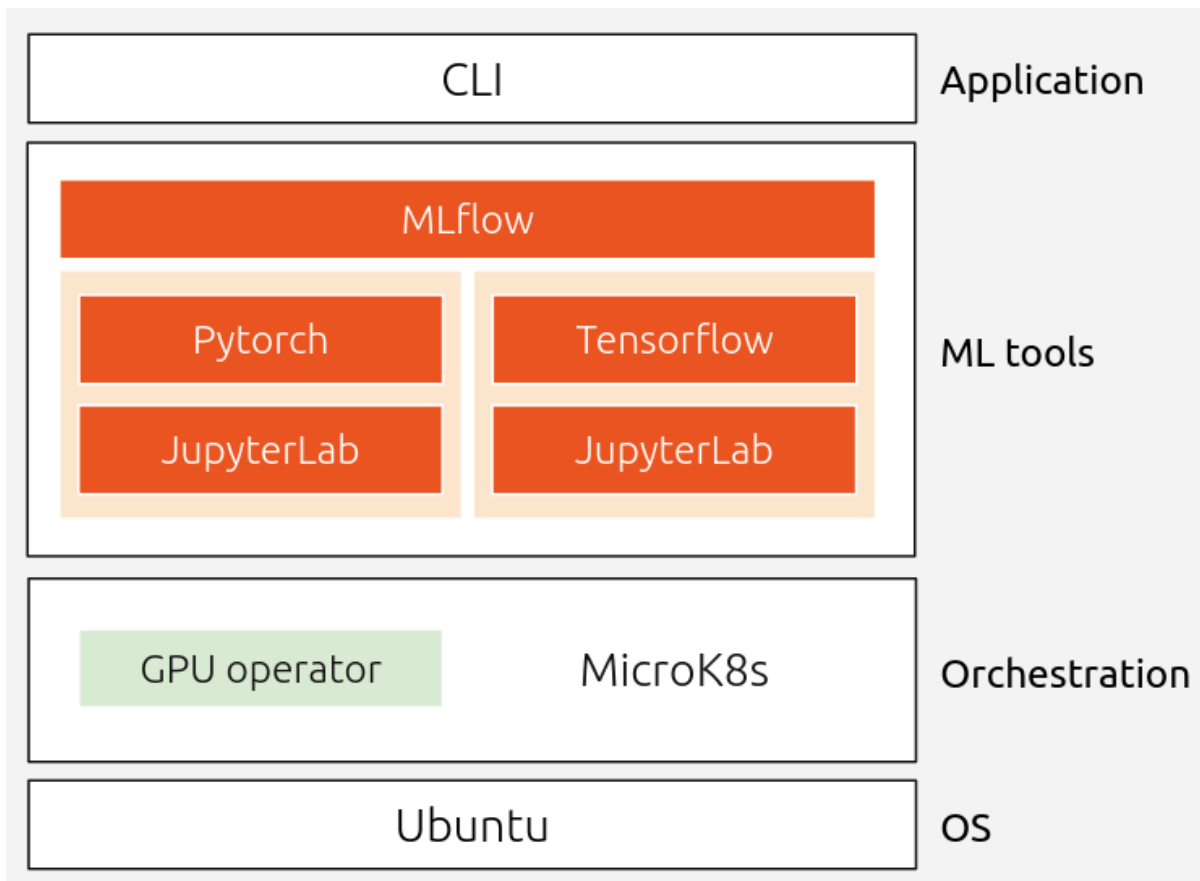


Fig. 1: Architecture overview

More details on each layer are discussed in the following sections.

### Application

DSS is a Command Line Interface (CLI)-based tool, accessible from the Ubuntu terminal. See [Manage DSS](#) to learn about how to manage your DSS environment and the available CLI commands.

### ML tools

DSS includes:

- Jupyter Notebooks: Open-source environment that provides a flexible interface to organise DS projects and ML workloads.
- MLflow: Open-source platform for managing the ML life cycle, including experiment tracking and model registry.
- ML frameworks: DSS comes by default with PyTorch and TensorFlow. Users can manually add other frameworks, depending on their needs and use cases.

### Jupyter Notebooks

A [Jupyter Notebook](#) is essentially a [Kubernetes deployment](#), also known as *Pod*, running a Docker image with Jupyter Lab and a dedicated ML framework, such as PyTorch or TensorFlow. For each Jupyter Notebook, DSS mounts a [Hostpath](#) directory-backed persistent volume to the data directory. All Jupyter Notebooks share the same persistent volume, allowing them to exchange data seamlessly. The full path to that persistent volume is `/home/joyan/shared`.

### MLflow

MLflow operates in [local mode](#), meaning that metadata and artefacts are, by default, stored in a local directory.

This local directory is backed by a persistent volume, mounted to a Hostpath directory of the MLflow Pod. The persistent volume can be found in the directory `/mlruns`.

### Orchestration

DSS requires a container orchestration solution. DSS relies on [Canonical K8s](#), a lightweight Kubernetes distribution.

Therefore, Canonical K8s needs to be deployed before installing DSS on the host machine. It must be configured with local storage support to handle persistent volumes used by DSS.

### GPU support

DSS can run with or without the use of GPUs. If needed, follow [NVIDIA GPU Operator](#) for deployment details.

DSS does not automatically install the tools and libraries required for running GPU workloads. It relies on Canonical K8s for the required operating-system drivers. It also depends on the chosen image, for example, CUDA when working with NVIDIA GPUs.

#### Caution

GPUs from other silicon vendors rather than NVIDIA can be configured. However, its functionality is not guaranteed.

### Storage

DSS expects a default [storage class](#) in the Kubernetes deployment, which is used to persist Jupyter Notebooks and MLflow artefacts. In Canonical K8s, a local storage class should be configured to provision Kubernetes' *PersistentVolumeClaims* (PVCs).

A shared PVC is used across all Jupyter Notebooks to share and persist data. MLflow also uses its dedicated PVC to store the logged artefacts. This is the DSS default storage configuration and cannot be altered.

This choice ensures that all storage is backed up on the host machine in the event of cluster restarts.



**Note**

By default, you can access the DSS storage anytime under your local directory `/var/snap/k8s/common/default-storage`.

The following diagram summarises the DSS storage:

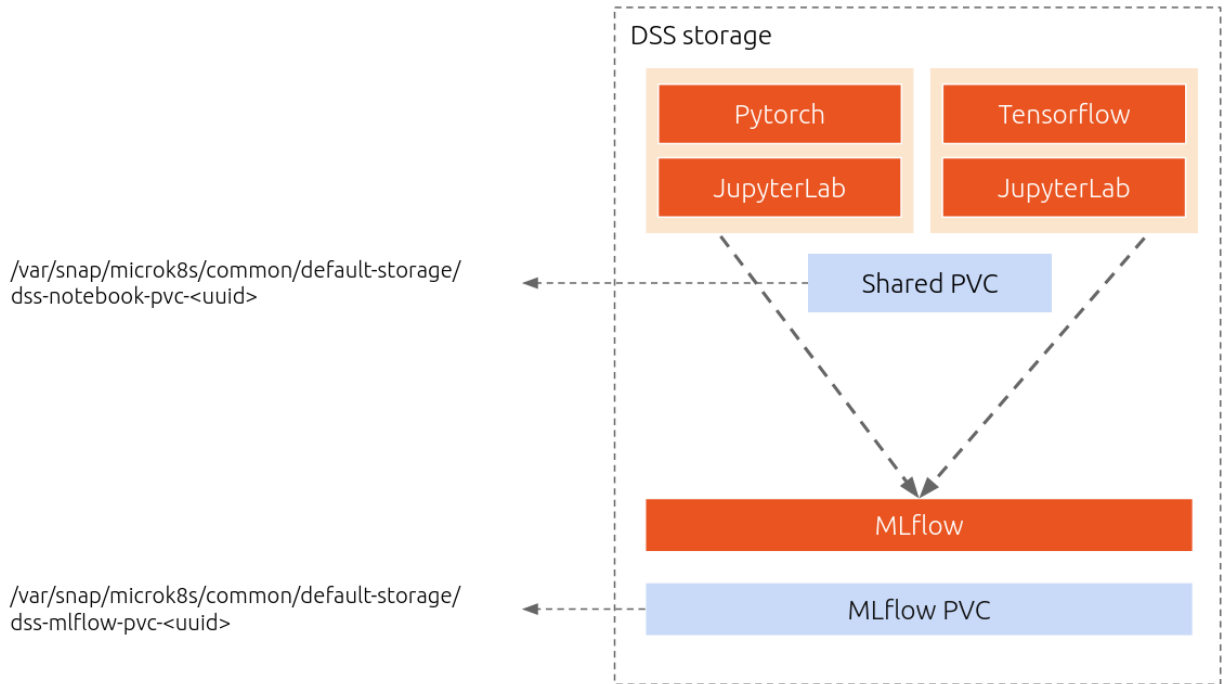


Fig. 2: Storage overview

### Operating system

DSS is native on Ubuntu, being developed, tested, and validated on it. Moreover, the solution can be used on any Linux distribution.

### Namespace configuration

DSS runs on a dedicated Kubernetes namespace. By default, it contains two Kubernetes Pods.

The NVIDIA GPU support runs on another dedicated namespace. This includes the GPU Operator for managing access and usage.

### Accessibility

Jupyter Notebooks and MLflow can be accessed from a web browser through the Pod IP that is given access through Canonical K8s. See [Access a notebook](#) and [Access MLflow](#) for more details.