

---

# Dask Yarn

Nov 08, 2018



---

# Contents

---

|          |  |          |
|----------|--|----------|
| <b>1</b> | <b>Install</b>                             | <b>3</b> |
| 1.1      | Quickstart . . . . .                       | 3        |
| 1.2      | Distributing Python Environments . . . . . | 6        |
| 1.3      | Configuration . . . . .                    | 7        |
| 1.4      | Submitting Applications . . . . .          | 8        |
| 1.5      | API Docs . . . . .                         | 9        |
| 1.6      | CLI Docs . . . . .                         | 11       |



Dask-Yarn deploys Dask on **YARN** clusters, such as are found in traditional Hadoop installations. Dask-Yarn provides an easy interface to quickly start, scale, and stop Dask clusters natively from Python.

```
from dask_yarn import YarnCluster
from dask.distributed import Client

# Create a cluster where each worker has two cores and eight GiB of memory
cluster = YarnCluster(environment='environment.tar.gz',
                      worker_vcores=2,
                      worker_memory="8GiB")
# Scale out to ten such workers
cluster.scale(10)

# Connect to the cluster
client = Client(cluster)
```

Dask-Yarn uses **Skein**, a Pythonic library to create and deploy YARN applications.



Dask-Yarn is designed to only require installation on an edge node. To install, use one of the following methods:

### Install with Conda:

```
conda install -c conda-forge dask-yarn
```

### Install with Pip:

```
pip install dask-yarn
```

### Install from Source:

Dask-Yarn is available on [github](#) and can always be installed from source.

```
pip install git+https://github.com/dask/dask-yarn.git
```

## 1.1 Quickstart

Dask-Yarn is designed to be used like any other python library - install it locally and use it in your code (either interactively, or as part of an application). As long as the computer you're deploying on has access to the YARN cluster (usually an edge node), everything should work fine.

### 1.1.1 Install Dask-Yarn on an Edge Node

Dask-Yarn is designed to be used from an edge node. To install, use either [conda](#) or [pip](#) to create a new environment and install `dask-yarn` on the edge node.

#### Conda Environments:

Create a new conda environment with `dask-yarn` installed. You may also want to add any other packages you rely on for your work.

```
$ conda create -n my_env dask-yarn      # Create an environment
$ conda activate my_env                 # Activate the environment
```

### Virtual Environments:

Create a new virtual environment with `dask-yarn` installed. You may also want to add any other packages you rely on for your work.

```
$ python -m venv my_env                 # Create an environment using venv
$ source my_env/bin/activate           # Activate the environment
$ pip install dask-yarn                 # Install some packages
```

## 1.1.2 Package your environment for Distribution

We need to ensure that the libraries used on the Yarn cluster are the same as what you are using locally. By default, `dask-yarn` handles this by distributing a packaged python environment to the Yarn cluster as part of the applications. This is typically handled using

- `conda-pack` for Conda environments
- `venv-pack` for virtual environments

See *Distributing Python Environments* for more information.

### Conda Environments:

If you haven't already installed `conda-pack`, you'll need to do so now. You can either install it in the environment to be packaged, or your root environment (where it will be available to use in all environments).

```
$ conda install -c conda-forge conda-pack # Install conda-pack
$ conda-pack                             # Package environment
Collecting packages...
Packing environment at '/home/username/miniconda/envs/my_env' to 'my_env.tar.gz'
[#####] | 100% Completed | 12.2s
```

### Virtual Environments:

If you haven't already installed `venv-pack`, you'll need to do so now.

```
$ pip install venv-pack                 # Install venv-pack
$ venv-pack                             # Package environment
Collecting packages...
Packing environment at '/home/username/my-env' to 'my-env.tar.gz'
[#####] | 100% Completed | 8.3s
```

## 1.1.3 Kinit (Optional)

If your cluster is configured to use Kerberos for authentication, you need to make sure you have an active ticket-granting-ticket before continuing:

```
$ kinit
```



### 1.1.4 Usage

To start a YARN cluster, create an instance of `YarnCluster`. This constructor takes several parameters, leave them empty to use the defaults defined in the `Configuration`.

```
from dask_yarn import YarnCluster
from dask.distributed import Client

# Create a cluster where each worker has two cores and eight GiB of memory
cluster = YarnCluster(environment='environment.tar.gz',
                      worker_vcores=2,
                      worker_memory="8GiB")

# Connect to the cluster
client = Client(cluster)
```

By default no workers are started on cluster creation. To change the number of workers, use the `YarnCluster.scale()` method. When scaling up, new workers will be requested from YARN. When scaling down, workers will be intelligently selected and scaled down gracefully, freeing up resources.

```
# Scale up to 10 workers
cluster.scale(10)

# ...

# Scale back down to 2 workers
cluster.scale(2)
```

Alternatively, if you're working interactively in a Jupyter Notebook or JupyterLab, you can use the provided graphical interface to change the cluster size.

In [4]: `cluster`

## YarnCluster

**Workers** 2

Workers

Scale

**Cores** 4

**Memory**8.58 GB

Normally the cluster will persist until the `YarnCluster` object is deleted. To be more explicit about when the cluster is shutdown, you can either use the cluster as a context manager, or manually call `YarnCluster.shutdown()`.

```
# Use ``YarnCluster`` as a context manager
with YarnCluster(...) as cluster:
    # The cluster will remain active inside this block,
    # and will be shutdown when the context exits.

# Or manually call `shutdown`
cluster = YarnCluster(...)
```

(continues on next page)

```
# ...
cluster.shutdown()
```

## 1.2 Distributing Python Environments

We need to ensure that the libraries used on the Yarn cluster are the same as what you are using locally. By default, `dask-yarn` handles this by distributing a packaged python environment to the Yarn cluster as part of the applications. This is typically handled using

- `conda-pack` for Conda environments
- `venv-pack` for virtual environments

These environments can contain any Python packages you might need, but require `dask-yarn` (and its dependencies) at a minimum.

### 1.2.1 Packing Conda Environments using Conda-Pack

You can package a conda environment using `conda-pack`.

```
$ conda create -n my-env dask-yarn scikit-learn          # Create an environment
$ conda activate my-env                                # Activate the environment
$ conda-pack                                           # Package environment
Collecting packages...
Packing environment at '/home/username/miniconda/envs/my-env' to 'my-env.tar.gz'
[#####] | 100% Completed | 12.2s
```

### 1.2.2 Packing Virtual Environments using Venv-Pack

You can package a virtual environment using `venv-pack`. The virtual environment can be created using either `venv` or `virtualenv`. Note that the python linked to in the virtual environment must exist and be accessible on every node in the YARN cluster. If the environment was created with a different Python, you can change the link path using the `--python-prefix` flag. For more information see the `venv-pack` documentation.

```
$ python -m venv my_env                                # Create an environment using venv
$ python -m virtualenv my_env                          # Or create an environment using
↳virtualenv
$ source my_env/bin/activate                           # Activate the environment
$ pip install dask-yarn scikit-learn                  # Install some packages
$ venv-pack                                           # Package environment
Collecting packages...
Packing environment at '/home/username/my-env' to 'my-env.tar.gz'
[#####] | 100% Completed | 8.3s
```

### 1.2.3 Specifying the Environment for the Cluster

You can now start a cluster with the packaged environment by passing the path to the constructor, e.g. `YarnCluster(environment='my-env.tar.gz', ...)`. After startup you may want to verify that your versions match with the following:

```
from dask_yarn import YarnCluster
from dask.distributed import Client

cluster = YarnCluster(environment='my-env.tar.gz')
client = Client(cluster)
client.get_versions(check=True) # check that versions match between all nodes
```

## 1.3 Configuration

Specifying all parameters to the `YarnCluster` constructor every time may be error prone, especially when sharing this workflow with new users. Alternatively, you can provide defaults in a configuration file, traditionally held in `~/config/dask/yarn.yaml` or `/etc/dask/yarn.yaml`. Note that this configuration is *optional*, and only changes the defaults when not specified in the constructor.

### Example:

```
# /home/username/.config/dask/yarn.yaml
yarn:
  name: dask # Application name
  queue: default # Yarn queue to deploy to

  environment: /path/to/my-env.tar.gz

  tags: [] # List of strings to tag applications

  scheduler: # Specifications of scheduler container
    vcores: 1
    memory: 4GiB

  worker: # Specifications of worker containers
    vcores: 2
    memory: 8GiB
    count: 0 # Number of workers to start on initialization
    restarts: -1 # Allowed number of restarts, -1 for unlimited
```

Users can now create `YarnClusters` without specifying any additional information.

```
from dask_yarn import YarnCluster

cluster = YarnCluster()
cluster.scale(20)
```

For more information on Dask configuration see the [Dask configuration documentation](#).

## 1.4 Submitting Applications

**Warning:** The submission API is experimental and may change between versions

Sometimes you have Dask Application you want to deploy completely on YARN, without having a corresponding process running on an edge node. This may come up with production applications deployed automatically, or long running jobs you don't want to consume edge node resources.

To handle these cases, `dask-yarn` provides a *CLI Docs* that can be used to submit applications to be run on the YARN cluster asynchronously. There are three commands that may be useful here:

- `dask-yarn submit`: submit an application to the YARN cluster
- `dask-yarn status`: check on the status of an application
- `dask-yarn kill`: kill a running application

### 1.4.1 Submitting an Application

To prepare an application to be submitted using `dask-yarn submit`, you need to change the creation of your `YarnCluster` from using the constructor to using `YarnCluster.from_current()`.

```
# Replace this
cluster = YarnCluster(...)

# with this
cluster = YarnCluster.from_current()
```

This is because the script won't be run until the cluster is already created - at that point configuration passed to the `YarnCluster` constructor won't be useful. Cluster configuration is instead passed via the `dask-yarn submit` CLI (note that as before, the cluster can be scaled dynamically after creation).

```
# Submit `myscript.py` to run on a dask cluster with 8 workers,
# each with 2 cores and 4 GiB
$ dask-yarn submit \
  --environment my_env.tar.gz \
  --worker-count 8 \
  --worker-vcores 2 \
  --worker-memory 4GiB \
  myscript.py

application_1538148161343_0051
```

This outputs a YARN Application ID, which can be used with other YARN tools.

### 1.4.2 Checking Application Status

Submitted application status can be checked using the YARN Web UI, or programmatically using `dask-yarn status`. This command takes one parameter - the application id.

```
$ dask-yarn status application_1538148161343_0051
APPLICATION_ID      NAME      STATE      STATUS      CONTAINERS
↪VCORES    MEMORY    RUNTIME
application_1538148161343_0051    dask    RUNNING    UNDEFINED    9
↪          33792    6m          17
```

### 1.4.3 Killing a Running Application

Submitted applications normally run until completion. If you need to terminate one before then, you can use the `dask-yarn kill` command. This command takes one parameter - the application id.

```
$ dask-yarn kill application_1538148161343_0051
```

### 1.4.4 Accessing the Application Logs

Application logs can be retrieved a few ways:

- The logs of running applications can be viewed using the [Skein Web UI](#) (dask-yarn is built using [Skein](#)).
- The logs of completed applications can be viewed using the `yarn logs` command.

```
$ yarn logs --applicationId application_1538148161343_0051
```

## 1.5 API Docs

```
class dask_yarn.YarnCluster (environment=None, n_workers=None, worker_vcores=None,
worker_memory=None, worker_restarts=None, worker_env=None,
scheduler_vcores=None, scheduler_memory=None, de-
ploy_mode=None, name=None, queue=None, tags=None,
skein_client=None)
```

Start a Dask cluster on YARN.

You can define default values for this in Dask's `yarn.yaml` configuration file. See <http://docs.dask.org/en/latest/configuration.html> for more information.

#### Parameters

**environment** [str, optional] Path to an archived Python environment (either `tar.gz` or `zip`).

**n\_workers** [int, optional] The number of workers to initially start.

**worker\_vcores** [int, optional] The number of virtual cores to allocate per worker.

**worker\_memory** [str, optional] The amount of memory to allocate per worker. Accepts a unit suffix (e.g. '2 GiB' or '4096 MiB'). Will be rounded up to the nearest MiB.

**worker\_restarts** [int, optional] The maximum number of worker restarts to allow before failing the application. Default is unlimited.

**worker\_env** [dict, optional] A mapping of environment variables to their values. These will be set in the worker containers before starting the dask workers.

**scheduler\_vcores** [int, optional] The number of virtual cores to allocate per scheduler.

**scheduler\_memory** [str, optional] The amount of memory to allocate to the scheduler. Accepts a unit suffix (e.g. '2 GiB' or '4096 MiB'). Will be rounded up to the nearest MiB.

**deploy\_mode** [{‘remote’, ‘local’}, optional] The deploy mode to use. If ‘remote’, the scheduler will be deployed in a YARN container. If ‘local’, the scheduler will run locally, which can be nice for debugging. Default is ‘remote’.

**name** [str, optional] The application name.

**queue** [str, optional] The queue to deploy to.

**tags** [sequence, optional] A set of strings to use as tags for this application.

**skein\_client** [skein.Client, optional] The `skein.Client` to use. If not provided, one will be started.

### Examples

```
>>> cluster = YarnCluster(environment='my-env.tar.gz', ...)
>>> cluster.scale(10)
```

**close** (\*\*kwargs)

Close this cluster. An alias for `shutdown`.

**See also:**

`shutdown`

**dashboard\_link**

Link to the dask dashboard. None if dashboard isn’t running

**classmethod from\_application\_id** (app\_id, skein\_client=None)

Connect to an existing `YarnCluster` with a given application id.

**Parameters**

**app\_id** [str] The existing cluster’s application id.

**skein\_client** [skein.Client] The `skein.Client` to use. If not provided, one will be started.

**Returns**

`YarnCluster`

**classmethod from\_current** ()

Connect to an existing `YarnCluster` from inside the cluster.

**Returns**

`YarnCluster`

**classmethod from\_specification** (spec, skein\_client=None)

Start a dask cluster from a skein specification.

**Parameters**

**spec** [skein.ApplicationSpec, dict, or filename] The application specification to use. Must define at least one service: ‘dask.worker’. If no ‘dask.scheduler’ service is defined, a scheduler will be started locally.

**skein\_client** [skein.Client, optional] The `skein.Client` to use. If not provided, one will be started.

**scale** (n)

Scale cluster to n workers.

**Parameters**

**n** [int] Target number of workers

**Examples**

```
>>> cluster.scale(10) # scale cluster to ten workers
```

**scale\_down** (*workers*)

Retire the selected workers.

**Parameters**

**workers: list** List of addresses of workers to close.

**scale\_up** (*n*, *workers=None*)

Ensure there are atleast *n* dask workers available for this cluster.

No-op if *n* is less than the current number of workers.

**Examples**

```
>>> cluster.scale_up(20) # ask for twenty workers
```

**shutdown** (*status='SUCCEEDED'*, *diagnostics=None*)

Shutdown the application.

**Parameters**

**status** [{*'SUCCEEDED'*, *'FAILED'*, *'KILLED'*}, optional] The yarn application exit status.

**diagnostics** [str, optional] The application exit message, usually used for diagnosing failures. Can be seen in the YARN Web UI for completed applications under “diagnostics”. If not provided, a default will be used.

**workers** ()

A list of all currently running worker containers.

## 1.6 CLI Docs

**Warning:** The CLI is experimental and may change between versions

### 1.6.1 dask-yarn

Deploy Dask on Apache YARN

```
usage: dask-yarn [--help] [--version] command ...
```

**--help, -h**

Show this help message then exit

## Dask Yarn

---

### **--version**

Show version then exit

### **dask-yarn kill**

Kill a Dask application

```
usage: dask-yarn kill [--help] APP_ID
```

### **app\_id**

The application id

### **--help, -h**

Show this help message then exit

### **dask-yarn services**

Manage Dask services

```
usage: dask-yarn services [--help] command ...
```

### **--help, -h**

Show this help message then exit

### **dask-yarn services client**

Start a Dask client process

```
usage: dask-yarn services client [--help] script [args...]
```

### **script**

Path to a Python script to run.

### **args**

Any additional arguments to forward to *script*

### **--help, -h**

Show this help message then exit

### **dask-yarn services scheduler**

Start a Dask scheduler process

```
usage: dask-yarn services scheduler [--help]
```

### **--help, -h**

Show this help message then exit

### **dask-yarn services worker**

Start a Dask worker process



```
usage: dask-yarn services worker [--nthreads NTHREADS]
                                [--memory_limit MEMORY_LIMIT] [--help]
```

**--nthreads** <nthreads>

Number of threads. Defaults to number of vcores in container

**--memory\_limit** <memory\_limit>

Maximum memory available to the worker. This can be an integer (in bytes), a string (like '5 GiB' or '500 MiB'), or 0 (no memory management). Defaults to the container memory limit.

**--help, -h**

Show this help message then exit

## dask-yarn status

Check the status of a submitted Dask application

```
usage: dask-yarn status [--help] APP_ID
```

**app\_id**

The application id

**--help, -h**

Show this help message then exit

## dask-yarn submit

Submit a Dask application to a YARN cluster

```
usage: dask-yarn submit [--name NAME] [--queue QUEUE] [--tags TAGS]
                        [--environment ENVIRONMENT]
                        [--deploy-mode DEPLOY_MODE]
                        [--worker-count WORKER_COUNT]
                        [--worker-vcores WORKER_VCORES]
                        [--worker-memory WORKER_MEMORY]
                        [--worker-restarts WORKER_RESTARTS]
                        [--worker-env WORKER_ENV]
                        [--client-vcores CLIENT_VCORES]
                        [--client-memory CLIENT_MEMORY]
                        [--client-env CLIENT_ENV]
                        [--scheduler-vcores SCHEDULER_VCORES]
                        [--scheduler-memory SCHEDULER_MEMORY] [--help]
script [args...]
```

**script**

Path to a python script to run on the client

**args**

Any additional arguments to forward to *script*

**--name** <name>

The application name

**--queue** <queue>

The queue to deploy to

- tags** <tags>  
A comma-separated list of strings to use as tags for this application.
- environment** <environment>  
Path to an archived Python environment (either `tar.gz` or `zip`).
- deploy-mode** <deploy\_mode>  
Either 'remote' (default) or 'local'. If 'remote', the scheduler and client will be deployed in a YARN container. If 'local', they will be run locally.
- worker-count** <worker\_count>  
The number of workers to initially start.
- worker-vcores** <worker\_vcores>  
The number of virtual cores to allocate per worker.
- worker-memory** <worker\_memory>  
The amount of memory to allocate per worker. Accepts a unit suffix (e.g. '2 GiB' or '4096 MiB'). Will be rounded up to the nearest MiB.
- worker-restarts** <worker\_restarts>  
The maximum number of worker restarts to allow before failing the application. Default is unlimited.
- worker-env** <worker\_env>  
Environment variables to set on the workers. Pass a key-value pair like `--worker-env key=val`. May be used more than once.
- client-vcores** <client\_vcores>  
The number of virtual cores to allocate for the client.
- client-memory** <client\_memory>  
The amount of memory to allocate for the client. Accepts a unit suffix (e.g. '2 GiB' or '4096 MiB'). Will be rounded up to the nearest MiB.
- client-env** <client\_env>  
Environment variables to set on the client. Pass a key-value pair like `--client-env key=val`. May be used more than once.
- scheduler-vcores** <scheduler\_vcores>  
The number of virtual cores to allocate for the scheduler.
- scheduler-memory** <scheduler\_memory>  
The amount of memory to allocate for the scheduler. Accepts a unit suffix (e.g. '2 GiB' or '4096 MiB'). Will be rounded up to the nearest MiB.
- help, -h**  
Show this help message then exit

## Symbols

- `-client-env <client_env>`  
dask-yarn-submit command line option, 14
- `-client-memory <client_memory>`  
dask-yarn-submit command line option, 14
- `-client-vccores <client_vccores>`  
dask-yarn-submit command line option, 14
- `-deploy-mode <deploy_mode>`  
dask-yarn-submit command line option, 14
- `-environment <environment>`  
dask-yarn-submit command line option, 14
- `-help, -h`  
dask-yarn command line option, 11  
dask-yarn-kill command line option, 12  
dask-yarn-services command line option, 12  
dask-yarn-services-client command line option, 12  
dask-yarn-services-scheduler command line option, 12  
dask-yarn-services-worker command line option, 13  
dask-yarn-status command line option, 13  
dask-yarn-submit command line option, 14
- `-memory_limit <memory_limit>`  
dask-yarn-services-worker command line option, 13
- `-name <name>`  
dask-yarn-submit command line option, 13
- `-nthreads <nthreads>`  
dask-yarn-services-worker command line option, 13
- `-queue <queue>`  
dask-yarn-submit command line option, 13
- `-scheduler-memory <scheduler_memory>`  
dask-yarn-submit command line option, 14
- `-scheduler-vccores <scheduler_vccores>`  
dask-yarn-submit command line option, 14
- `-tags <tags>`  
dask-yarn-submit command line option, 13
- `-version`  
dask-yarn command line option, 11
- `-worker-count <worker_count>`  
dask-yarn-submit command line option, 14

- `-worker-env <worker_env>`  
dask-yarn-submit command line option, 14
- `-worker-memory <worker_memory>`  
dask-yarn-submit command line option, 14
- `-worker-restarts <worker_restarts>`  
dask-yarn-submit command line option, 14
- `-worker-vccores <worker_vccores>`  
dask-yarn-submit command line option, 14

## A

- `app_id`  
dask-yarn-kill command line option, 12  
dask-yarn-status command line option, 13
- `args`  
dask-yarn-services-client command line option, 12  
dask-yarn-submit command line option, 13

## C

- `close()` (dask\_yarn.YarnCluster method), 10

## D

- `dashboard_link` (dask\_yarn.YarnCluster attribute), 10
- `dask-yarn` command line option  
-help, -h, 11  
-version, 11
- `dask-yarn-kill` command line option  
-help, -h, 12  
app\_id, 12
- `dask-yarn-services` command line option  
-help, -h, 12
- `dask-yarn-services-client` command line option  
-help, -h, 12  
args, 12  
script, 12
- `dask-yarn-services-scheduler` command line option  
-help, -h, 12
- `dask-yarn-services-worker` command line option  
-help, -h, 13  
-memory\_limit <memory\_limit>, 13

`-nthreads <nthreads>`, 13

`dask-yarn-status` command line option

`-help, -h`, 13

`app_id`, 13

`dask-yarn-submit` command line option

`-client-env <client_env>`, 14

`-client-memory <client_memory>`, 14

`-client-vcores <client_vcores>`, 14

`-deploy-mode <deploy_mode>`, 14

`-environment <environment>`, 14

`-help, -h`, 14

`-name <name>`, 13

`-queue <queue>`, 13

`-scheduler-memory <scheduler_memory>`, 14

`-scheduler-vcores <scheduler_vcores>`, 14

`-tags <tags>`, 13

`-worker-count <worker_count>`, 14

`-worker-env <worker_env>`, 14

`-worker-memory <worker_memory>`, 14

`-worker-restarts <worker_restarts>`, 14

`-worker-vcores <worker_vcores>`, 14

`args`, 13

`script`, 13

## F

`from_application_id()` (`dask_yarn.YarnCluster` class method), 10

`from_current()` (`dask_yarn.YarnCluster` class method), 10

`from_specification()` (`dask_yarn.YarnCluster` class method), 10

## S

`scale()` (`dask_yarn.YarnCluster` method), 10

`scale_down()` (`dask_yarn.YarnCluster` method), 11

`scale_up()` (`dask_yarn.YarnCluster` method), 11

`script`

`dask-yarn-services-client` command line option, 12

`dask-yarn-submit` command line option, 13

`shutdown()` (`dask_yarn.YarnCluster` method), 11

## W

`workers()` (`dask_yarn.YarnCluster` method), 11

## Y

`YarnCluster` (class in `dask_yarn`), 9