Bacalhau FAQs

Can I use multiple data sources in the same job?

You can use the `--input` or `-i` flag multiple times with multiple different CIDs, URLs or S3 objects, and give each of them a path to be mounted at.

For example, doing `bacalhau run cat/main.wasm -i ipfs://CID1:/input1 -i ipfs://CID2:/input2` will result in both the `input1` and `input2` folders being available to your running WASM with the CID contents. You can use `-i` as many times as you need.

How can I submit Job requests through CLI to communicate with my Node directly?

To ensure the CLI can communicate with our node directly (`bacalhau --api-host <MY_NODE_PUBLIC_IP> version`), you need to make sure the **1234** port is open.

Why does my API server listening on `/ip4/127.0.0.1/tcp/5001` when I invoke IPFS Daemon when fetching an IPFS Multiaddress?

Bacalhau communicates with IPFS via it's API port and not the swarm port which is why it's **5001** and not **4001**.

The key thing is whether the IPFS node is running on the same host as the Bacalhau daemon. If it is, then **127.0.0.1** is enough to route traffic between the two (because they are both on the same node). If IPFS is running on a different node than Bacalhau, then we need to replace **127.0.0.1** with the IP that the IPFS server is running on.

What to do when I get error connection refused when running Bacalhau API?

Problem

When running `bacalhau --api-host <MY_NODE_PUBLIC_IP> version` and you get this error message:

```
Error running version: publicapi: after posting request: Post
"http://127.0.0.1:1234/version": dial tcp 127.0.0.1:1234: connect:
```

```
connection refused
```

What to do

First, you'll need to check that the bacalhau server is up and running on the same host then it should be connecting using `127.0.0.1`. This can be checked by running `telnet 127.0.0.1 1234`. If telnet is not connecting to **127.0.0.1 1234** on the machine that bacalhau is running then one of 3 things:

- Bacalhau is running on a different machine
- it's running on a different port
- it's not running

Can I run Bacalhau in a containerized setup (nested containers)?

We don't support this as it will result in the classic Dind(Docker In Docker) problem.

Can I run non Docker jobs?

Yes! You can run programs using WebAssembly instead. See the onboarding WebAssembly for information on how to do that.

How do I run a script that requires installing packages from a package repository like pypi or apt?

Networking is supported by Bacalhau which enables one to run a script that requires packages from external repository. This is only for Docker workloads

How do I see a job's progress while it's running?

That's currently not possible

How do I get an IPFS peer if I want to start Bacalhau Server?

A viable option is to run your own IPFS daemon and fetch your multiaddress as explained here.

How can I download and query SQLite when it complains about being in read-only directory?

When downloading content to run your code against, it is written to a read-only directory. Unfortunately, by default, SQLite requires the directory to be writable so that it can create utility files during its use.

If you run your command with the `immutable` setting set to 1, then it will work. From the sqlite3 command you can use `.open 'file:/inputs/database.db?immutable=1'` where you should replace "database.db" with your downloaded database filename.

Can I run bacalhau serve on my home machine? What are the requirements?

You can run `bacalhau serve` on any machine that fits the prerequisites listed here.

> 💡 **TIP**
>
> The walkthrough in the docs has been tested only on Ubuntu 22, bacalhau is being developed on Linux/macOS environments and therefore should work fine there as well. However, Windows hosts are supported with limitations.

Can I stop a running job?

Yes. Given a valid job ID, you can use the cancel command to cancel the job, and stop it from running.

✏️ Edit this page

Last updated on **Sep 14, 2023**