

Computing in SHiP production cluster

Alexander Baranov, Konstantin Nikitin, Andrey Ustyuzhanin


Skygrid

System to run computations in docker containers:


- 1700 cores (60 machines)

- Data backends: dCache, afs, cvmfs, git

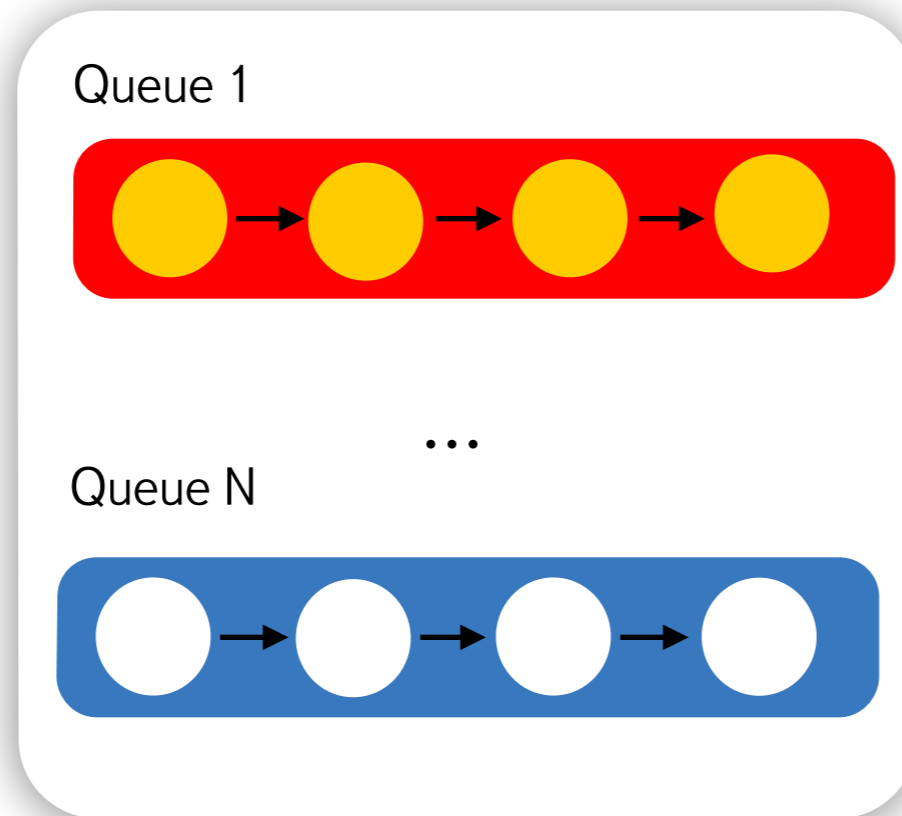
- Deployed in Yandex, can be extended to multiple sites



Users






Services



Metascheduler

(queue service)



Cluster



Cluster



Cluster

Workers

Done with skygrid

Already:

- | 1 billion muon background simulation
- | Neutrino and anti-neutrino background + reprocessing
- | 10^{10} muon background events

Plans:

- | Shield geometry optimisation
- | Pattern recognition for tracks
- | Event reconstruction

How computations are performed?

You start with your code and data

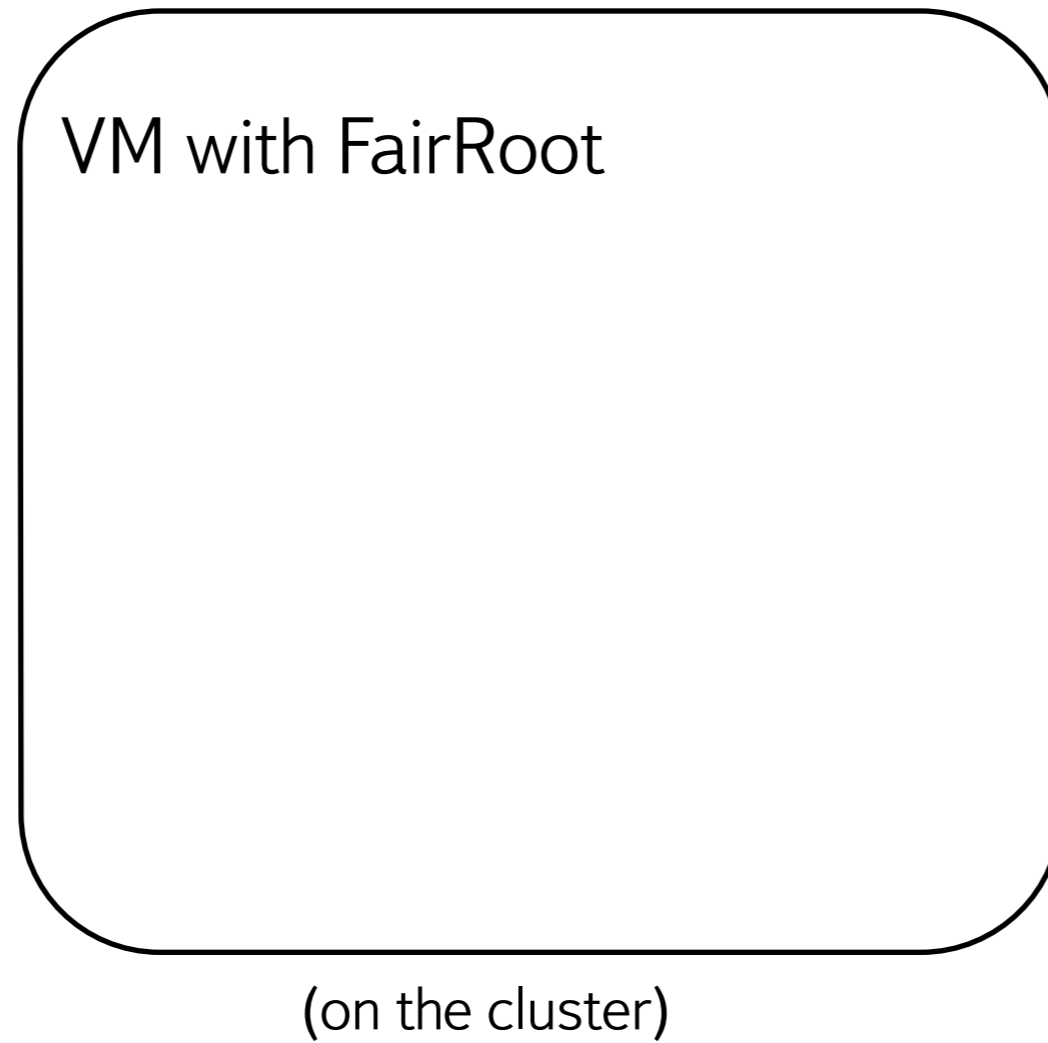
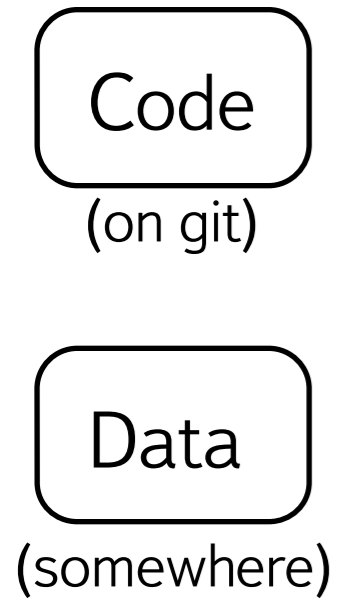
Code

(on git)

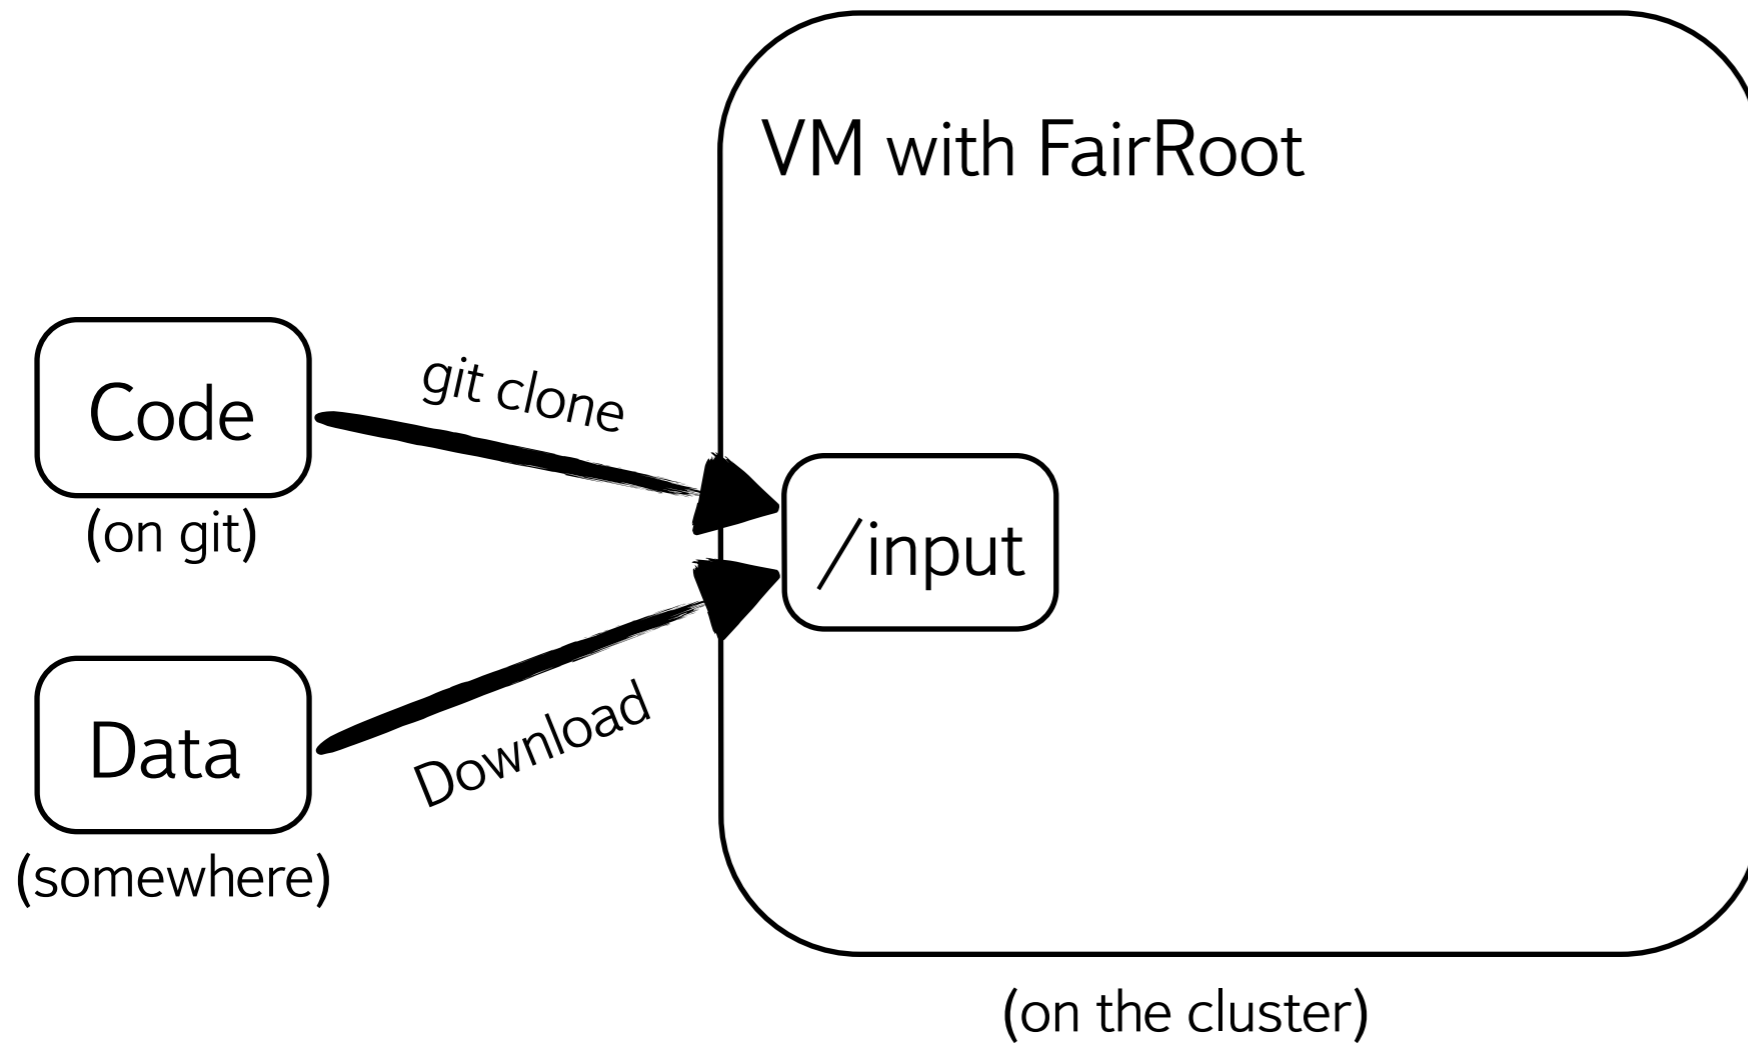
Data

(somewhere)

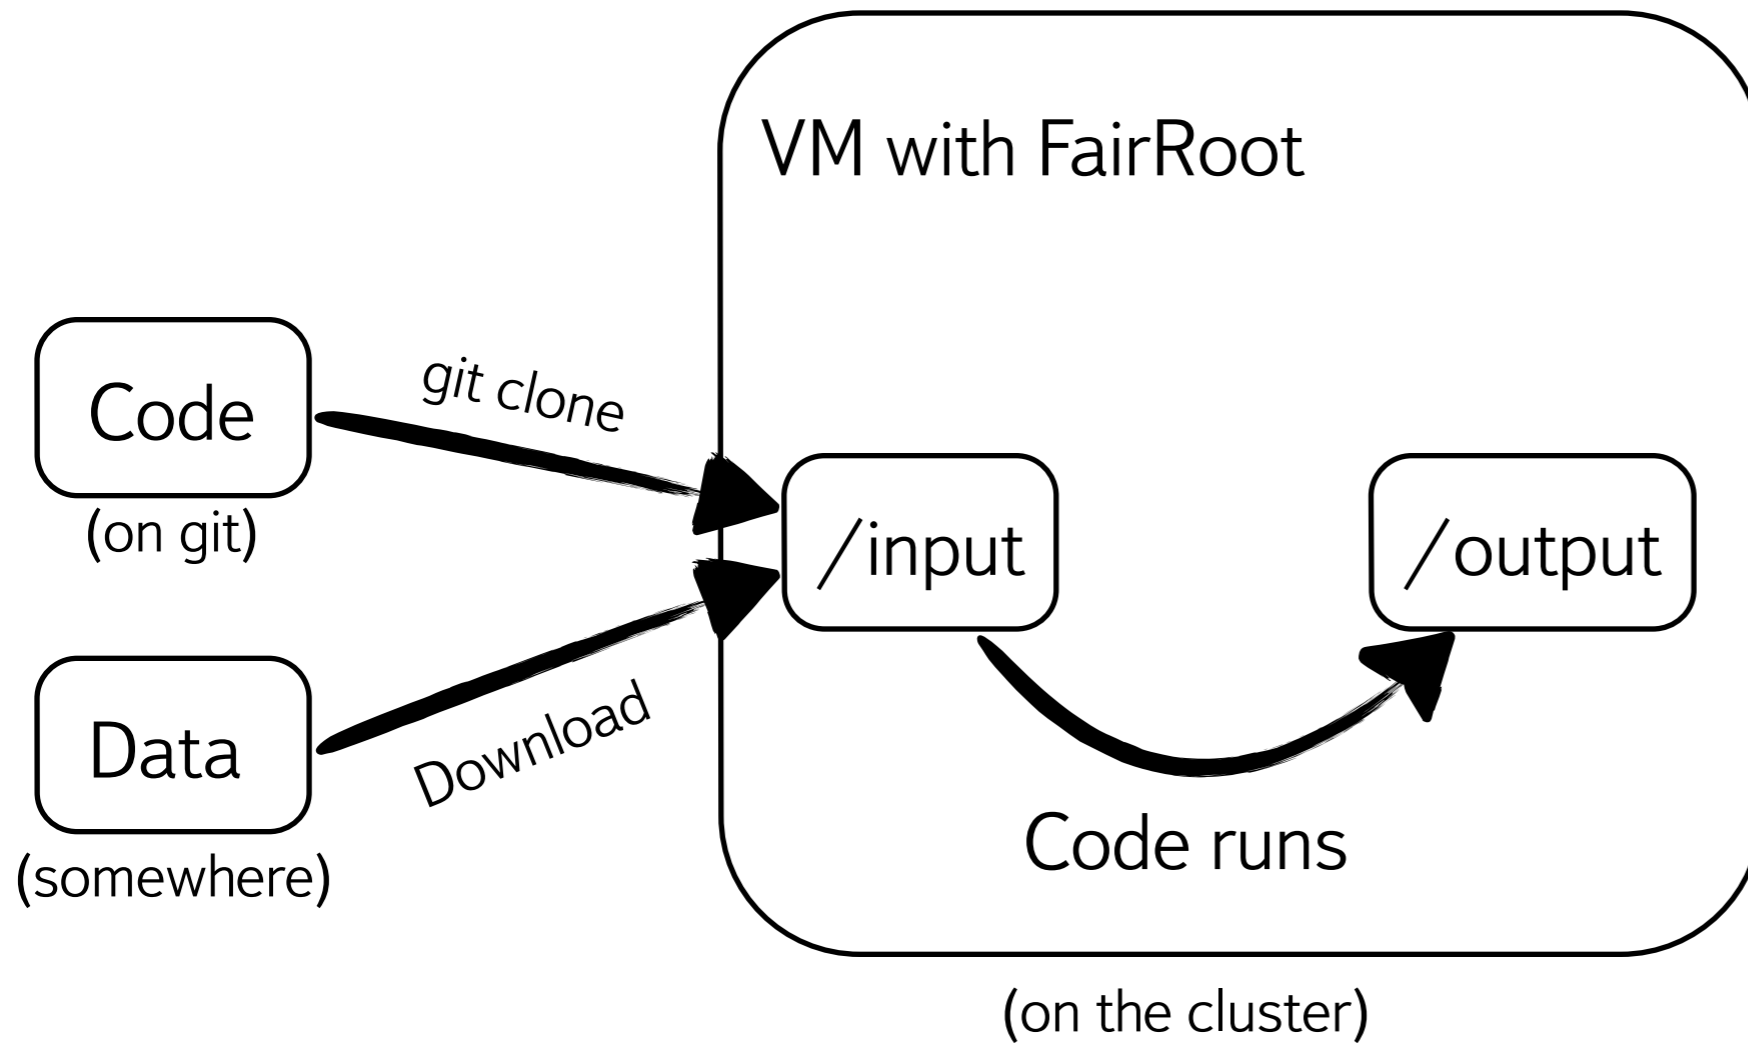
VM is started in the cluster



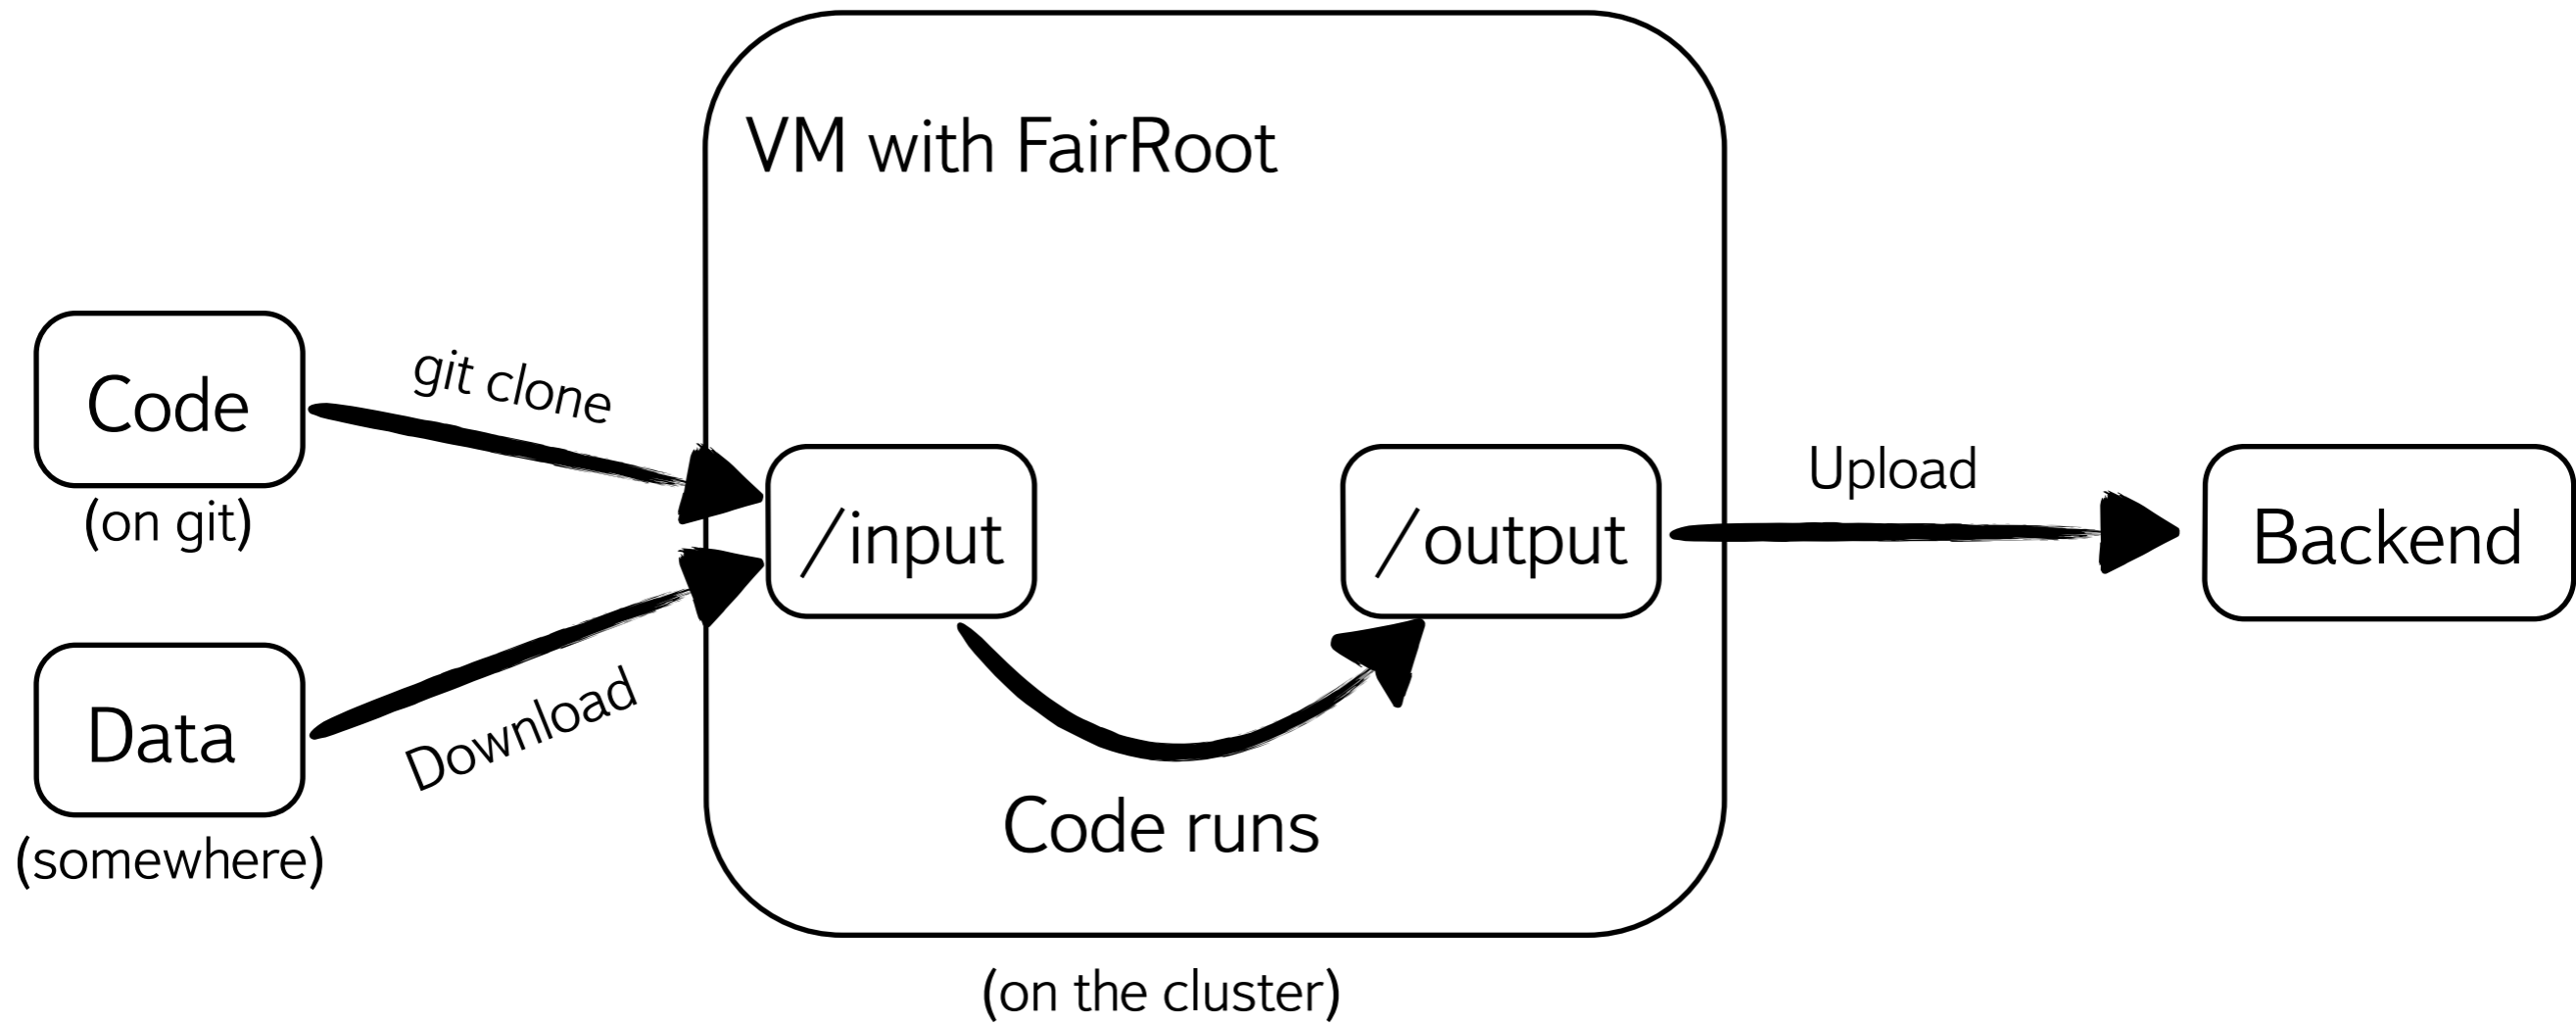
Your code / data is downloaded inside VM



Code runs and writes to /output



/output contents are stored



How do I submit jobs?

Check out tutorial: github.com/ShipSoft/FairShip/wiki/Docker-tutorial

To run your jobs we'll need next info:

- | Git repository with your code (should be public)
- | Command line to execute
- | Optional: Which data is needed?
- | Optional: any non-default(FairRoot) containers to run

You can send me or on skygrid-users@cern.ch

Wrap-up

| Lots of things done with skygrid already

| **Your jobs are really welcome!**

| Commit your code to github and you are ready to run it in skygrid

| Bookkeeping is coming



Questions and Answers

Mailing list: skygrid-users@cern.ch

All your comments and jobs are welcome!

Backup