

Luca Visentin, Luca Munaron and Federico Alessandro Ruffinatti
Department of Life Sciences and Systems Biology, Università di Torino, Torino, Italy

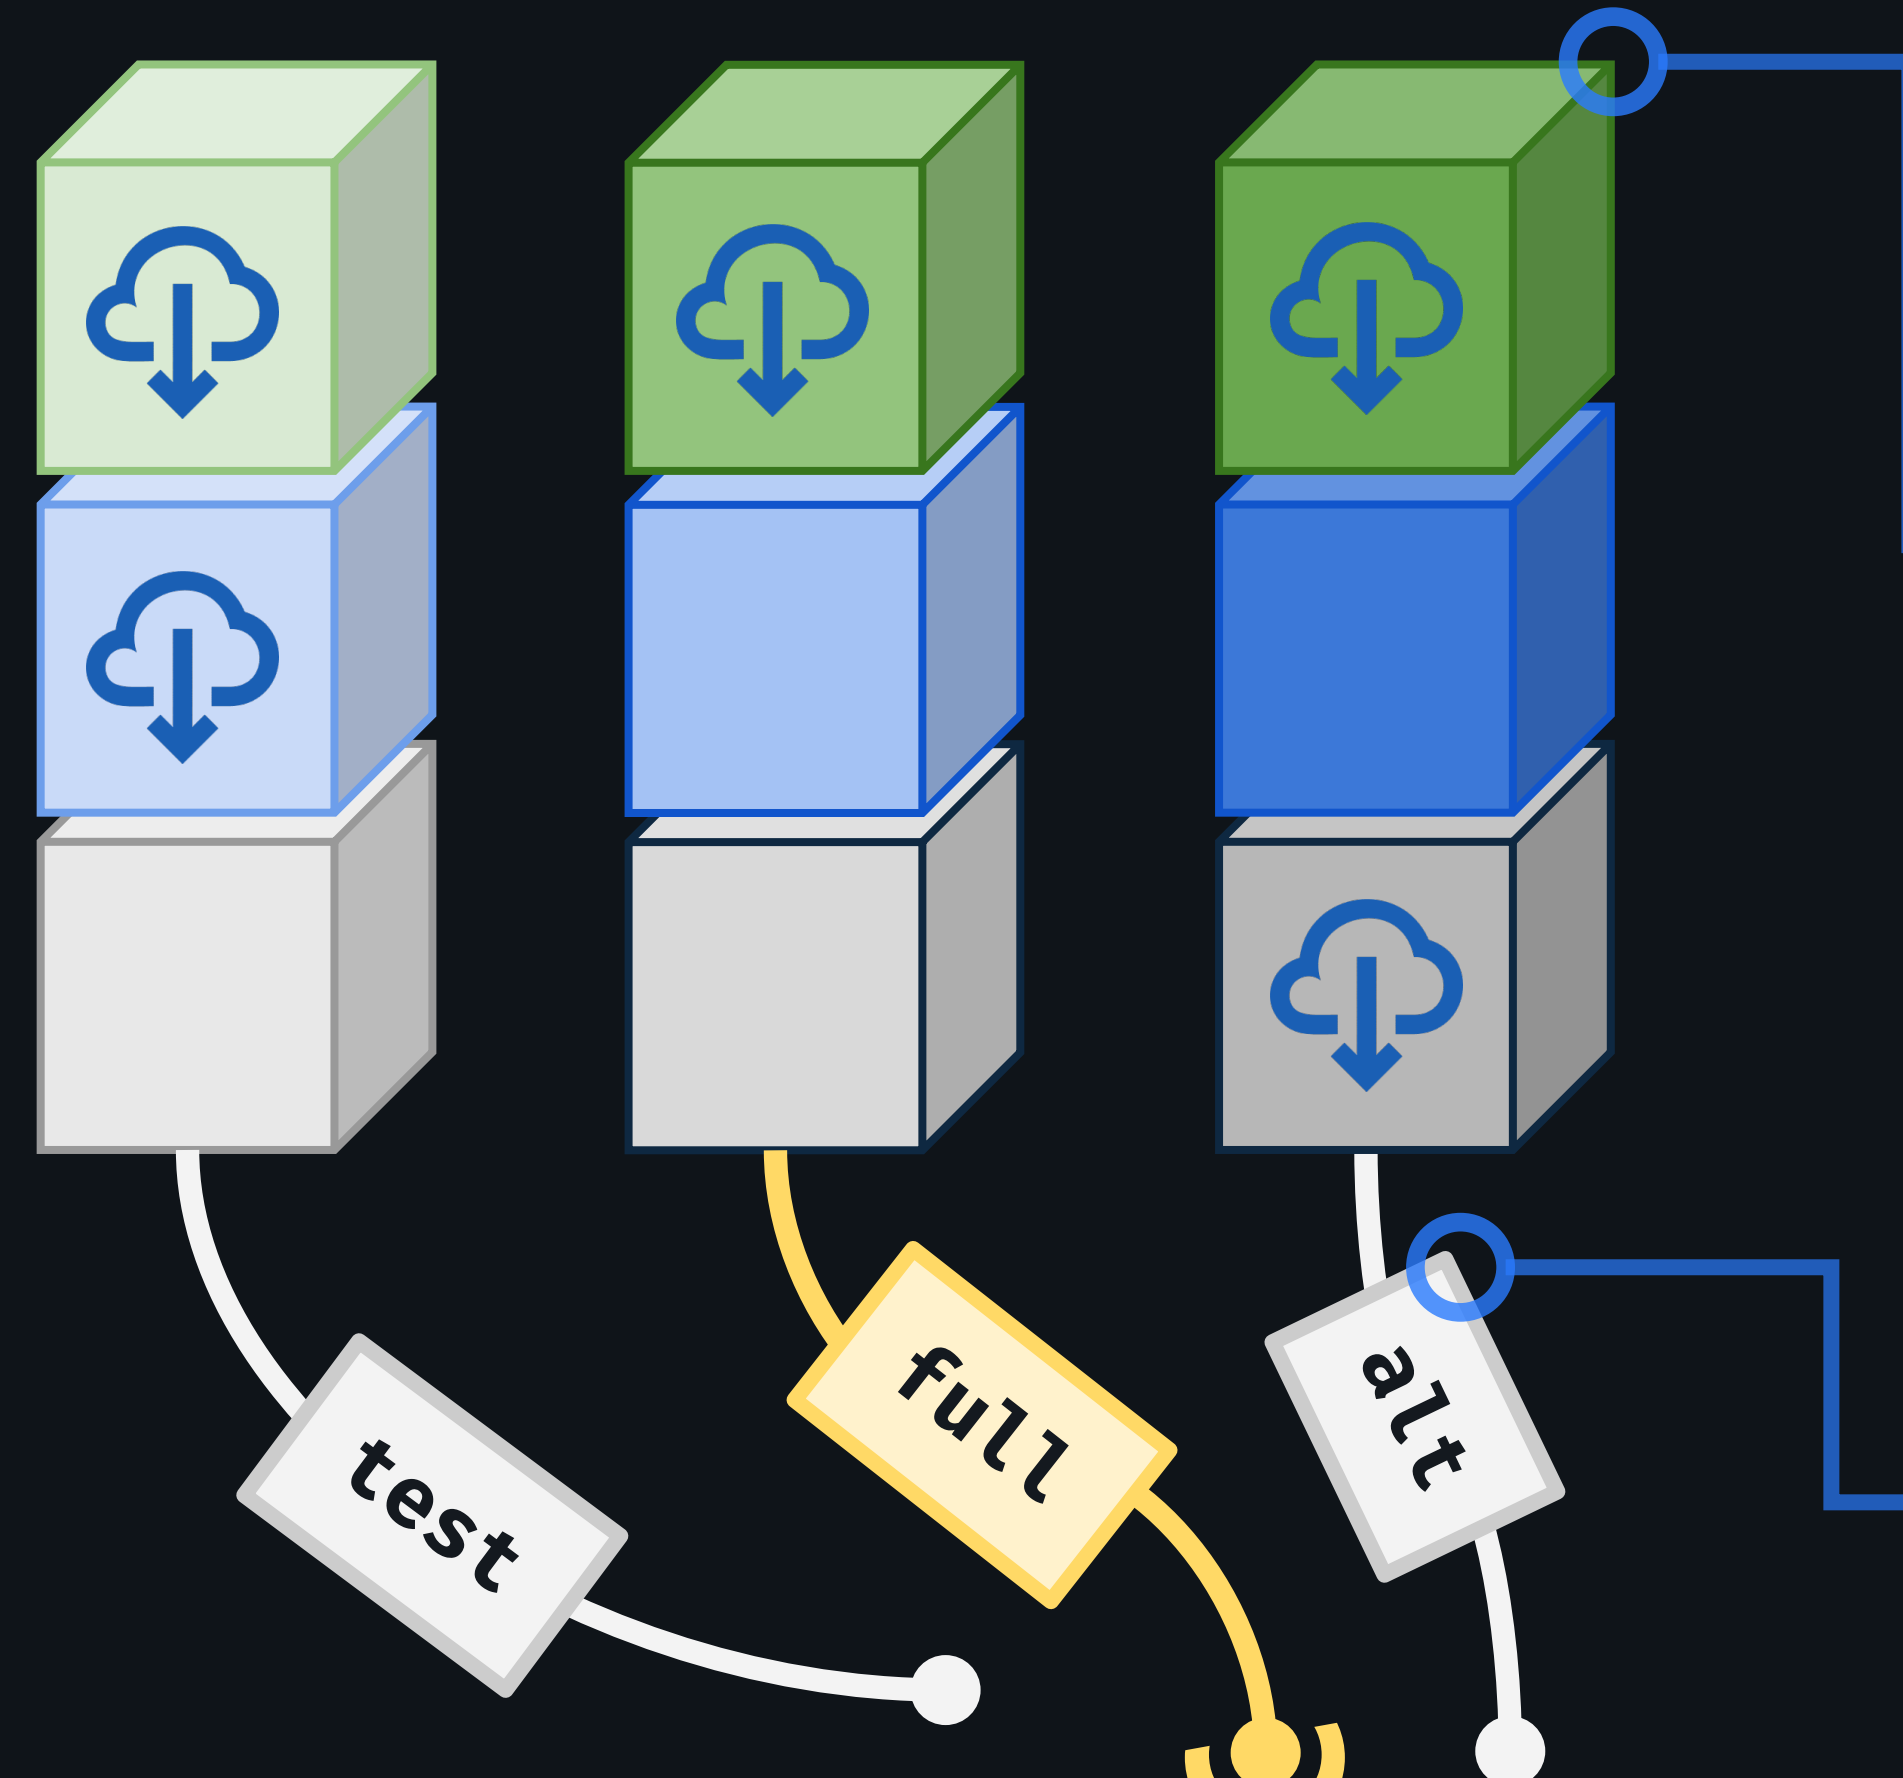
Quickstart.md

Kerblam! is a simple and lightweight project management system that helps you structure, execute, package and share your routine data analyses. It promotes a default project structure for ease of interoperability and collaboration.

With *Kerblam!* you can:

- setup new projects with sensible defaults;
- work reproducibly in containers without writing long, hard to remember commands;
- keep your workflows neat and ordered;
- bake-in documentation for your pipelines;
- manage your local data, including downloading, cleaning up of temporary and intermediate data, and packaging output;
- and more!

If you want data analysis, *Kerblam!* it!



kerblam data fetch

The push for FAIR(-er) data makes online datasets increasingly common. *Kerblam!* can fetch remote data for you, simply by configuring your endpoints. When you finish working on the project, you can also use *Kerblam!* to smartly delete these data and save precious disk space.

The declarative nature of the configuration makes it easy for others to re-run your project without having to ask you for input data.

Execution profiles

You often need (or want) to run the same workflows on different data: for testing purposes, on distinct batches of samples, or simply to change the configuration of your pipelines a bit.

Kerblam! allows you to do all of this declaratively and at runtime by creating *execution profiles* that map files to their alter-egos, so that *Kerblam!* can swap them at runtime for you on demand.

kerblam run

Kerblam! runs your pipelines, but not by itself: it leverages other workflow managers. It supports GNU make out-of-the-box, and everything else through shell scripts. This allows you to be flexible in the choice of workflow manager: are you familiar with *snakemake*? Use that! *nextflow*? Great! Do you just need a small Bash script? We love it!

Kerblam! can manage any number of pipelines, all with potentially different workflow managers, container images, documentation...

Easy containerisation

Kerblam! can use containers for you, so you don't have to remember how to *docker* run or write long volume mountpoints for your data if you want run inside containers what you normally would run locally (like *make*). Just write a *dockerfile* with the same name as your pipeline, and *Kerblam!* does the rest.

kerblam data clean

Kerblam! can clean your intermediate, temporary and even remotely-available input data with just one command, to keep your projects tidy and compact on disk.

Kerblam! can also show you how much disk space your data is using, broken down by file type: input, output and intermediate or temporary. With this, you can check whether your project is growing too large.

kerblam data pack

When you have to share your output with the rest of the world, *Kerblam!* can export your output data to a compressed tarball by carefully picking it out of the rest.

No more having to look up how to use *tar*!

kerblam.toml

Kerblam! uses the *kerblam.toml* file for its configuration, allowing you to work declaratively.

It's used to specify, among other things:


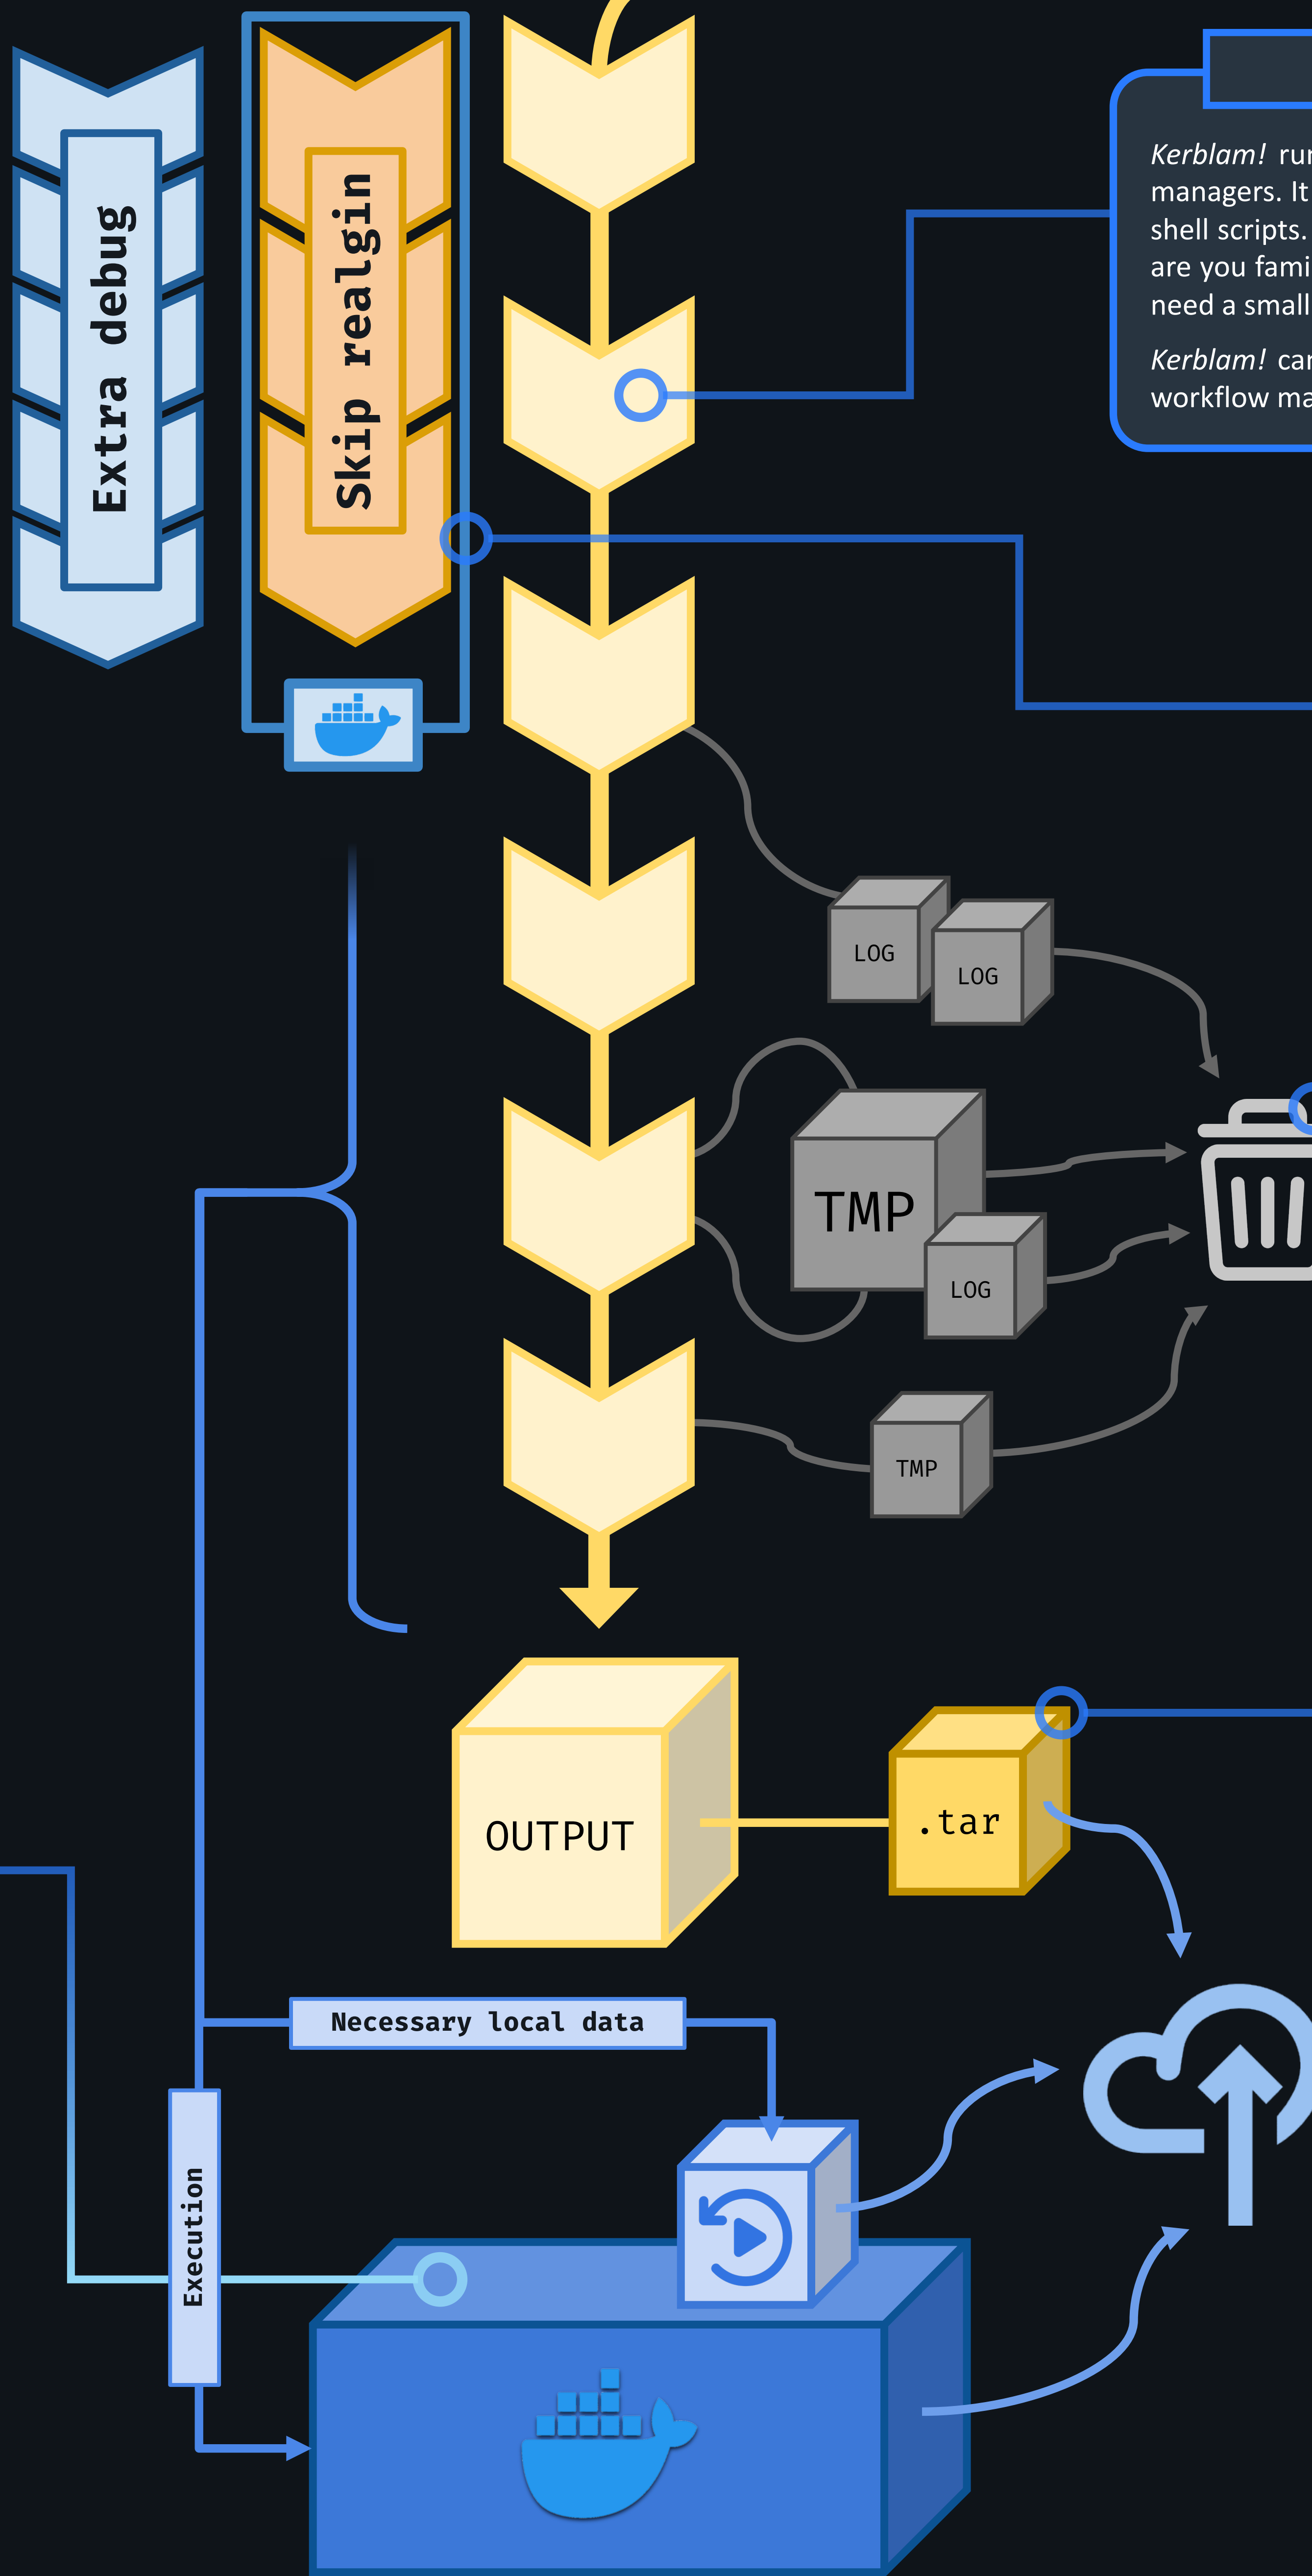
- what the folder structure of the project is;
- what files should be downloaded;
- which data profiles are available;
- what kind of backend execution you want to use: *Docker* or *Podman*;

kerblam package

When you are done, you usually want to share your work, so that others can reproduce your results.

Kerblam! allows you to package your pipelines into the same container images that you used for your own runs, export them for upload to a registry, and package the precious local-only data that others need for execution in a compressed *replay package*... just with one command!

If you share this replay package along with the container image, others can replay your work manually, or simply by using *kerblam replay*.




Read the docs

Hosted on



GitHub Pages

Docs generated by mdBook



Copyright: Luca Visentin, Luca Munaron, Federico Alessandro Ruffinatti, 2024/2025
This poster is CC-BY and available in Zenodo: <https://doi.org/10.5281/zenodo.11442700>

What's a Kerblam anyway? This project is named after the fictitious online shop/delivery company in S11E07 of Doctor Who. *Kerblam!* might be referred to as *Kerblam!*, *Kerblam* or *Kerblam*, interchangeably, although *Kerblam!* is preferred. The *Kerblam!* logo is written in the *kwark* font by *tup wanders*.

Written in Rust



With Love

Precompiled for



Source on GitHub

