



EBRAINS



Co-funded by  
the European Union

## ESD: Build Paths



Eric Müller

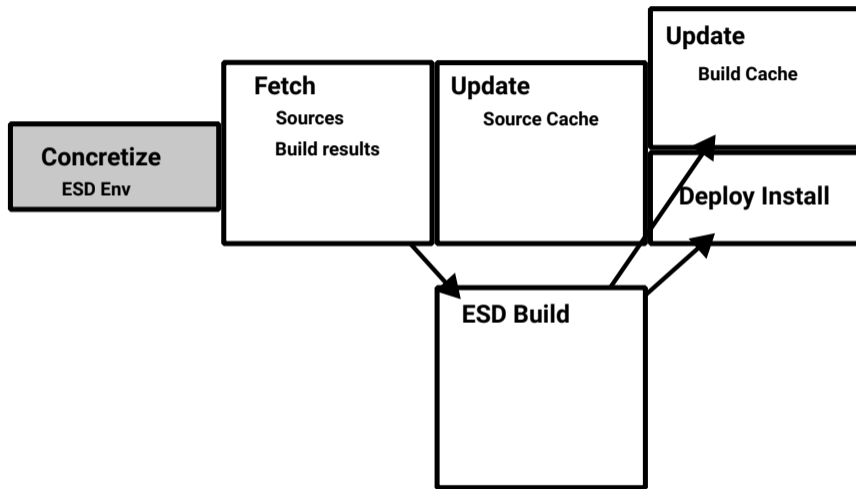
2024-11-28



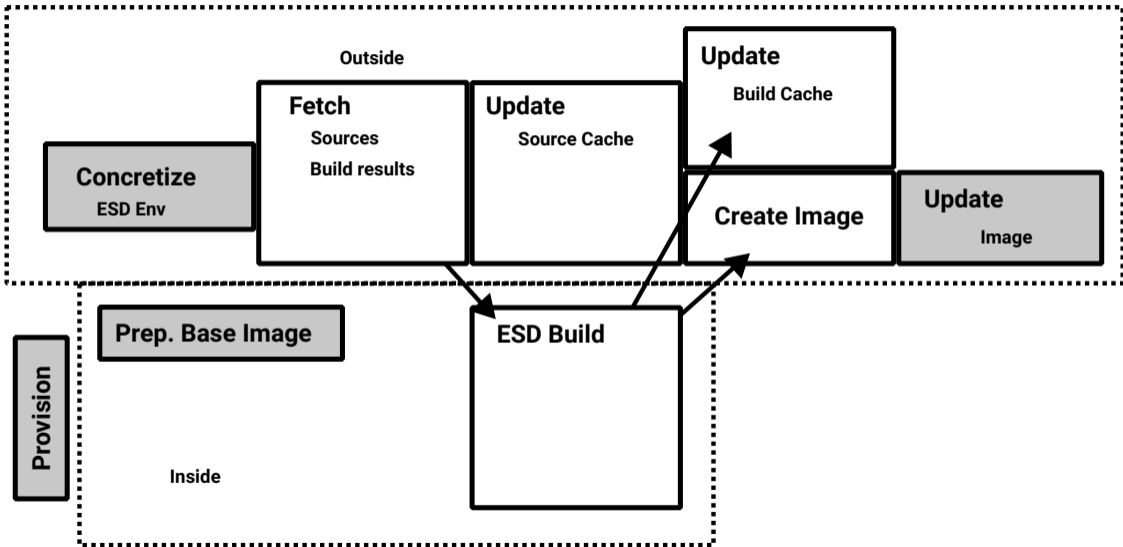
# Motivation

- Usage of the ESD...
  - in the lab,
  - on your laptop,
  - on HPC,
  - to run a (your) service.
- implies different build paths depending on
  - “target”-specific optional components (e.g., lab)
  - optimization
  - dependencies on underlying base install/image

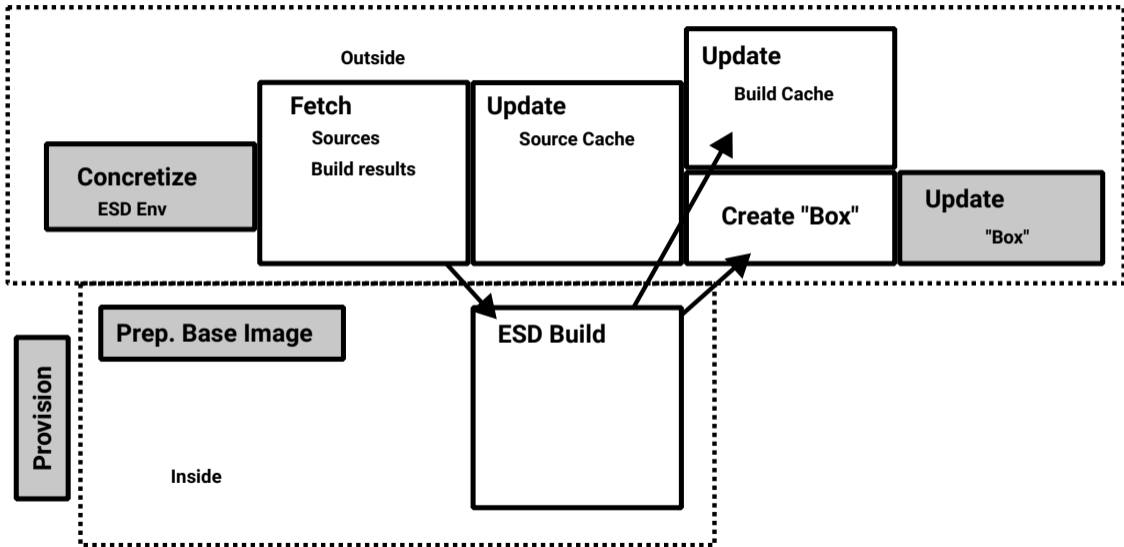
# (Full) Build Process for the Lab




# (Full) Build Process for Container Images



# (Full) Build Process for Virtual Machines



# (Full) Build Process for HPC Optimized Images




# Cacheable Entities and Build Persistence

- spack
  - sources
  - build results / `spack install buildcache`
- Images:
  - Base image / base image contents?
  - Final image
- Filesystem deployments
  - r/w distributed FS
  - separate write/deployment from r/o mount



# (Full) Build Process for HPC Optimized Images





# “Outside” & “Inside” Environments

- gitlab runners
  - OCI runtime-based container
  - could be different. . .
- image build envs:
  - “chroot”
  - container runtime
  - (thick) VM
- HPC builds:
  - online and offline build resources
- user CLI



# Encapsulation Levels

## “Pure” Userspace

(Type “0”)

- `mount --bind`
- `SECCOMP` or `ptrace()`
- (Type I requires `mount namespace`)

## User namespace

(Type II)

- “unprivileged”<sup>a</sup> or “rootless”
- requires user namespaces (`unshare(int)`)

---

<sup>a</sup>fails on unprivileged gitlab docker runners

## Classic OCI

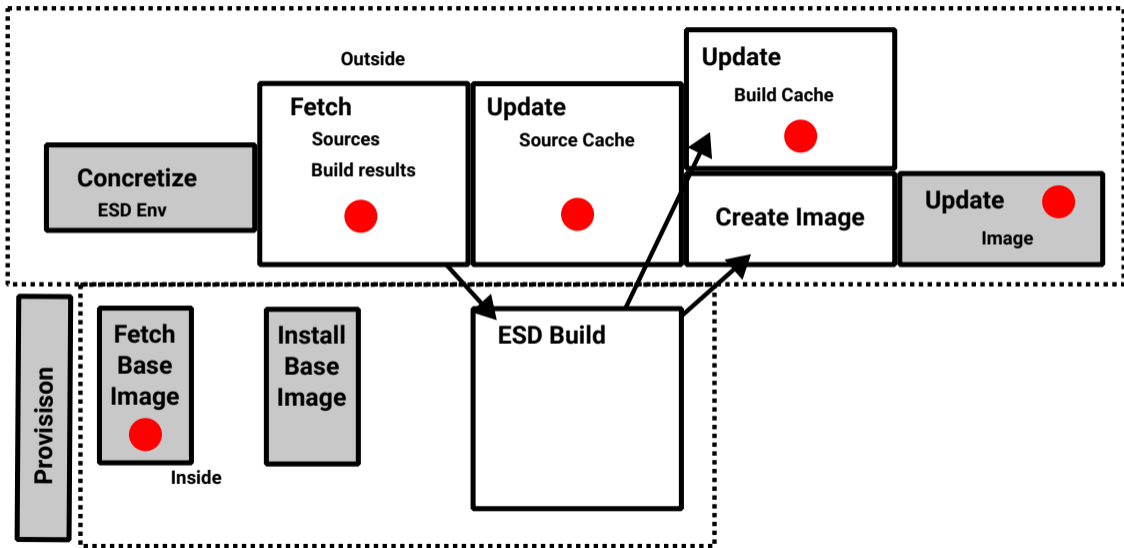
(Type III)

- not needed for software builds
- syscall-requiring package managers (`chown()`)...
- useful for certain service deployments

“Pure Userspace” is basically a Type 0 that doesn’t require `mount namespace` “Zero-consistency root emulation for unprivileged container image build”, Priedhorsky et al. [arXiv:2405.06085](https://arxiv.org/abs/2405.06085) [cs.DC].



# (Full) Build Process for HPC Optimized Images





# Additional Tools

- Getting “inside”
  - proot — Type 0
  - charliecloud — Type I
  - apptainer — Type II
  - buildah — Type II
- image file and OCI registry things
  - oras — handling artifacts in OCI registries
  - skopeo — handling OCI images and image repos
- PoC for modular/image ESD builds:  
<https://gitlab.ebrains.eu/ri/tech-hub/platform/esd/yashchiki>



# Misc Aspects

- Build path parallelism
  - spack env depfile
  - local, CI, scalable resources (unicore/pyslurm?)
  - per-package build requirements (e.g., memory per core)
- Filesystem recommendations for building software and images
  - `/dev/shm` recommended for server-type machines
  - avoid distributed filesystems (GPFS et al., NFS also not too much fun)
  - FUSE-mounted ones → depends (overlayfs etc. is fine)



# Summary

- Modular concept from the definition of the ESD (set of toplevel packages to a deployment)
- PoC implementation for image builds via “yashchiki”<sup>1</sup>

## Next steps:

- CI flow of (laptop) image operation on development branch
- modularization of build “paths” to run all mentioned combinations in CI
- Hands-on: Exploring container image builds on HPC/JUSUF

---

<sup>1</sup><https://gitlab.ebrains.eu/ri/tech-hub/platform/esd/yashchiki>