

The eHive system for pipeline management and execution

Part 1: System overview

Instructors:

Leo Gordon and Brandon Walts



<http://training.ensembl.org/events/2017/2017-03-23-ehiveRoslin>

Course overview

Day 1	Introduction to eHive
	Initializing and running eHive pipelines
	Pipeline configuration: modifying existing pipelines, and writing your own pipelines (part 1)
Day 2	Pipeline configuration (part 2)
	Writing your own runnable modules

Audience background

How many of you have experience:

- Running a pipeline or workflow
 - “By hand”
 - With a shell script
 - With a pipeline management tool
(Galaxy, Taverna, CWL, JobTree...)
- Submitting jobs to a compute cluster
(using LSF, SGE)
 - Using job control features of the scheduler
(job arrays, wait-for rules)
- Writing SQL queries
- Coding in Perl

eHive: an overview and brief history

At Ensembl, we run lots of workflows.

These pipelines have been getting more complex and running longer as the amount of data expands

We also have workflows that may not be complicated, but need to run in an automated way

In 2004, eHive was created to provide a standard, robust way to create and run analysis pipelines

Design goals:

- Scalable
- Adaptable
- Reproducible
- Traceable

eHive provides:

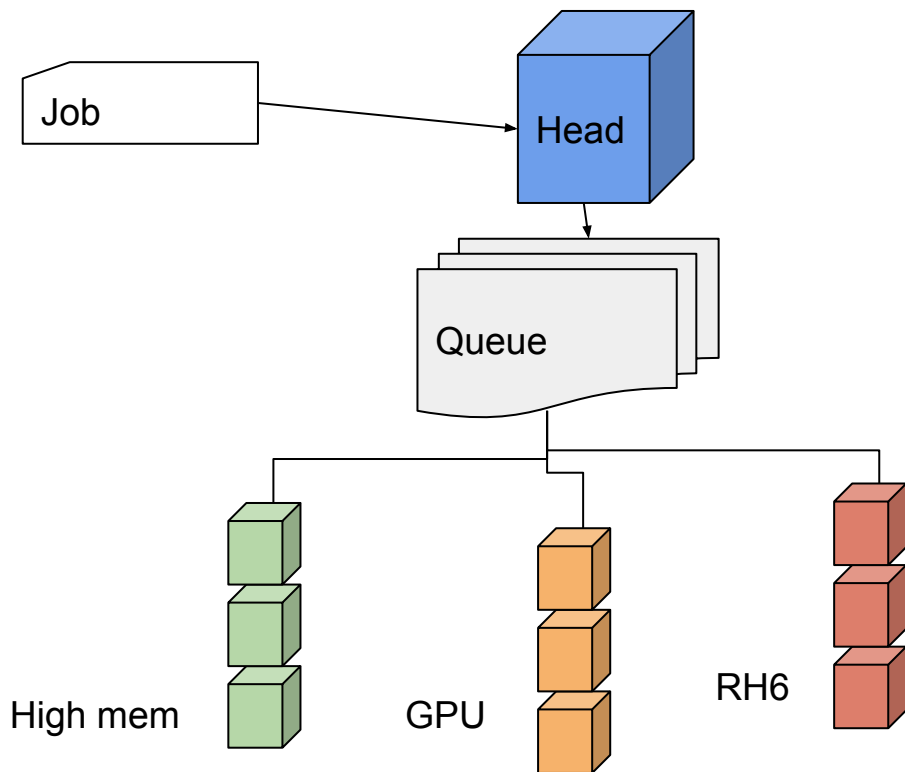
A framework to define workflows

- Sequencing operations
- Defining dependencies
- Flowing data
- Assigning resources

A system for executing workflows


- Job management
- Logging
- Checkpointing

Compute farms



- User at the head node submits job into a queue
- Scheduler assigns job onto a compute node based on
 - User priority
 - Resource requirements
 - Load
- More accurate load estimate = better chance to be scheduled
- Overhead for scheduling a job
- Some basic ability to handle relationships between jobs
 - Arrays
 - Wait-for

eHive's place in the farm



- eHive does not replace the scheduler (LSF, SGE, SLURM, etc.).
- It works one layer above the scheduler
- It is an automated, intelligent job **submission** engine
- It is an automated, intelligent job **execution** engine

eHive fundamentals

eHive's model of distributed computation

- Independent agents perform computation
- Coordination via a central list of jobs
- Agents are responsible for claiming jobs, performing work, and updating the central list

eHive fundamentals: walkthrough


Let's start with an analogy, building a house



image: David Wright, cc license

eHive fundamentals: walkthrough

What do we need to get our task done:



Blueprints: an overall plan. This plan is fixed and changes rarely if at all over the course of the project.



Job list: the current state of what has been done and what needs to be done next. This is being updated continuously as the project progresses.



Supervisor



Workers



Resources



Procedures: exactly how to perform a job

eHive fundamentals: walkthrough

The difference between the blueprints, the job list, and procedures



Blueprints are:

- Fixed throughout the whole project (in general)
- Describe “classes” of jobs
- Describe dependencies between job classes
- Describe resources required for jobs

Job lists are:

- Dynamic, jobs are being added throughout the project
- Describe “instances” of jobs
- Describe dependencies between individual jobs

eHive fundamentals: walkthrough

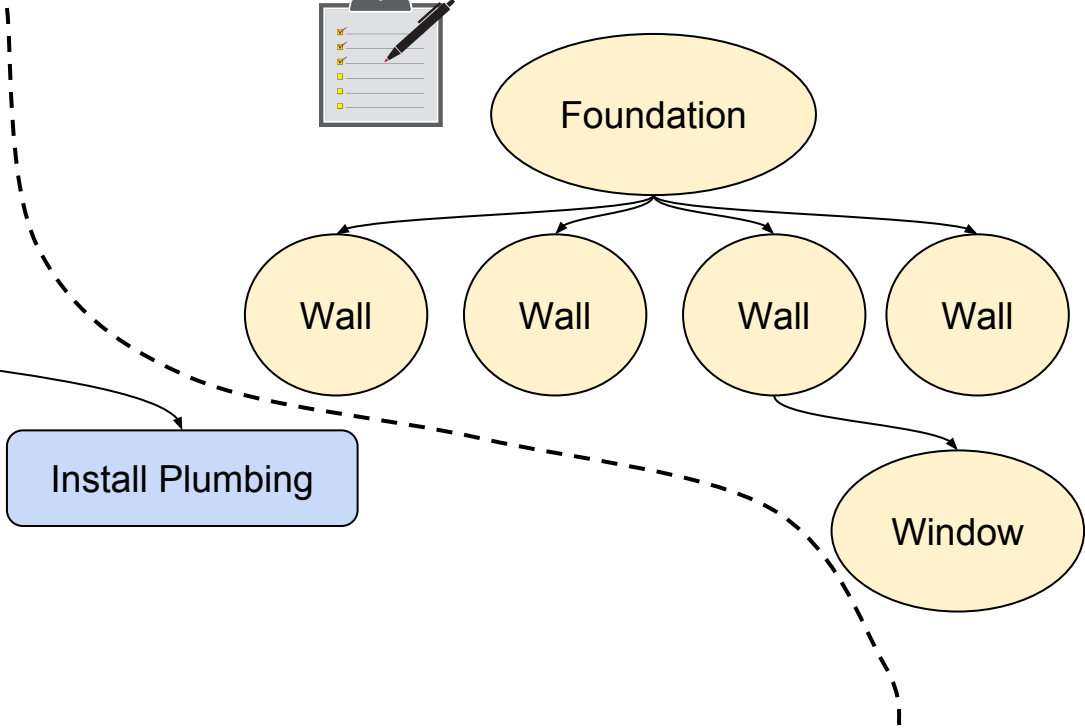
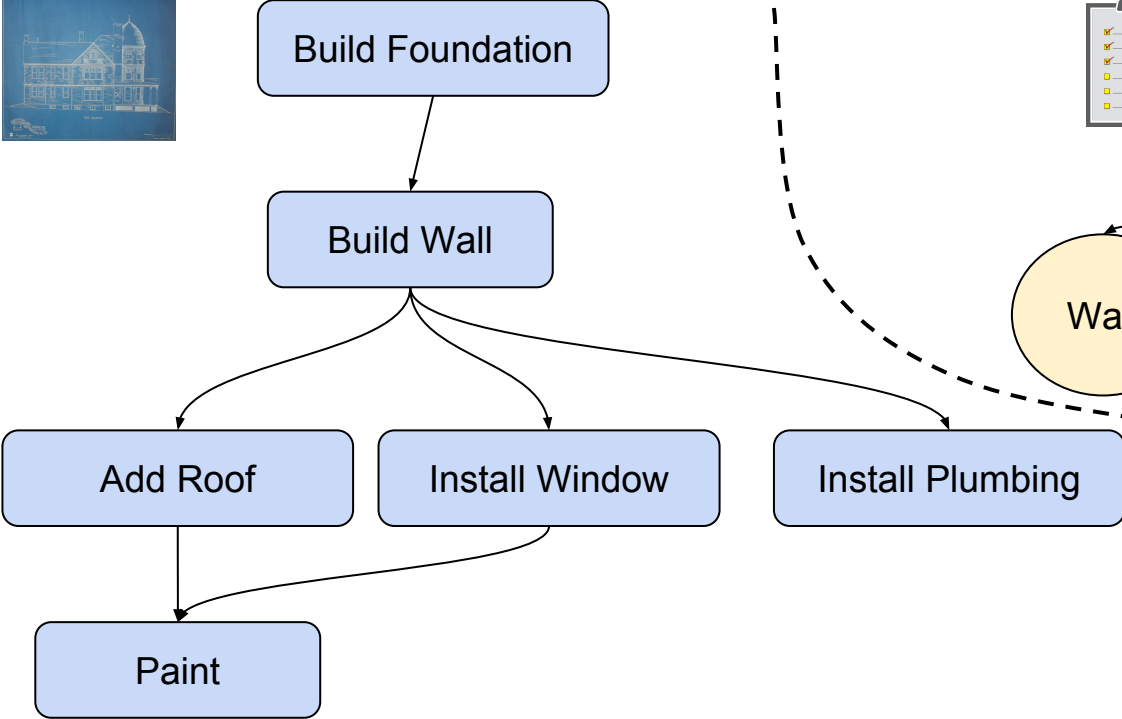
The difference between the blueprints, the job list, and procedures



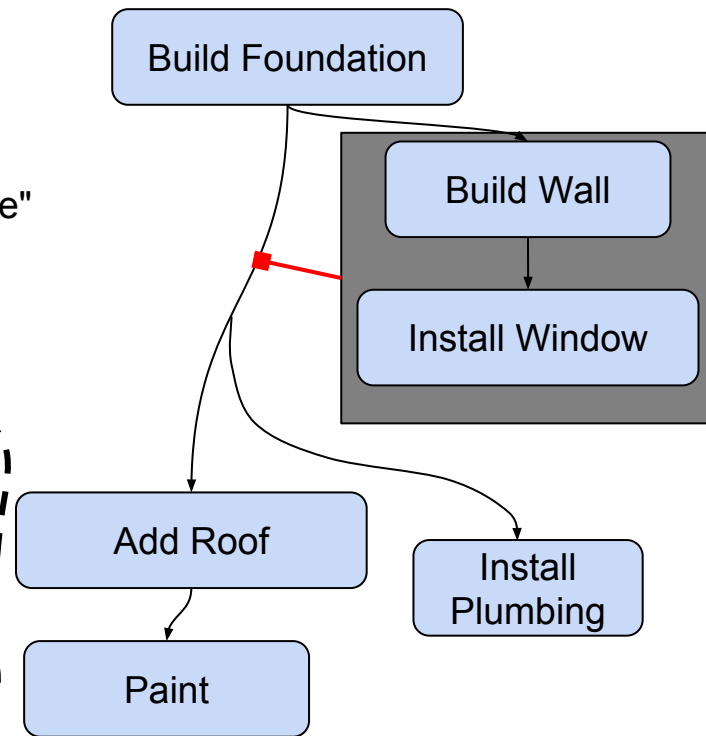
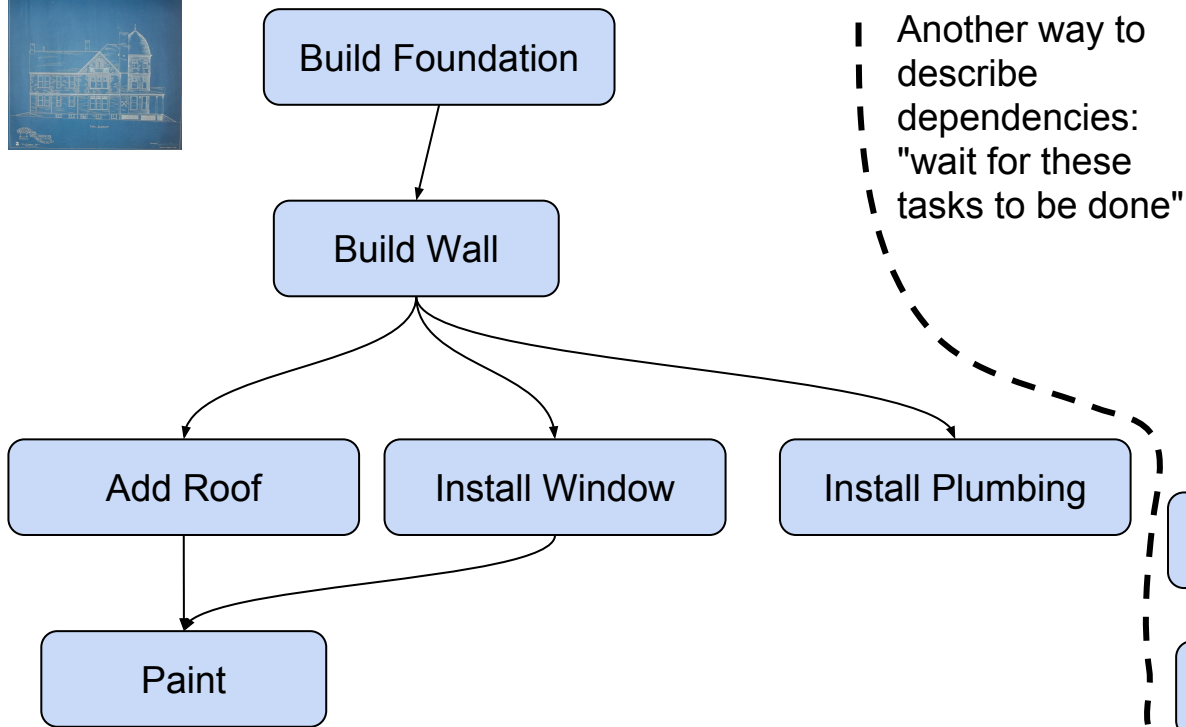
Procedures are:

- Fixed throughout the project
- Describe step by step how to do a particular job
- Are parameterized, so they can be customized by plugging in different parameters

eHive fundamentals: walkthrough



eHive fundamentals: walkthrough




eHive fundamentals: walkthrough

In eHive, the supervisor is very “hands off” and the workers take care of managing the jobs themselves.

Let's follow a worker to see how this happens...

eHive fundamentals: walkthrough

follow a worker...



1. Worker checks the job list and finds a job it can do



2. Worker claims a job, and specializes to perform it

eHive fundamentals: walkthrough

follow a worker...



3. Worker finishes job and marks it done on the job list



4. **Worker** checks the blueprints to see if it needs to add any new jobs to the list

eHive fundamentals: walkthrough

follow a worker...



5. Worker checks to see if there are any more jobs of the same type to do




6. If so, worker claims another job and starts to work on it




7. If not, worker checks to see if there are other types of jobs it can do

eHive fundamentals: walkthrough

follow a worker...



8. If so, worker claims the new job and respecializes to perform it



9. If not, worker leaves

eHive fundamentals: walkthrough

Worker time limit



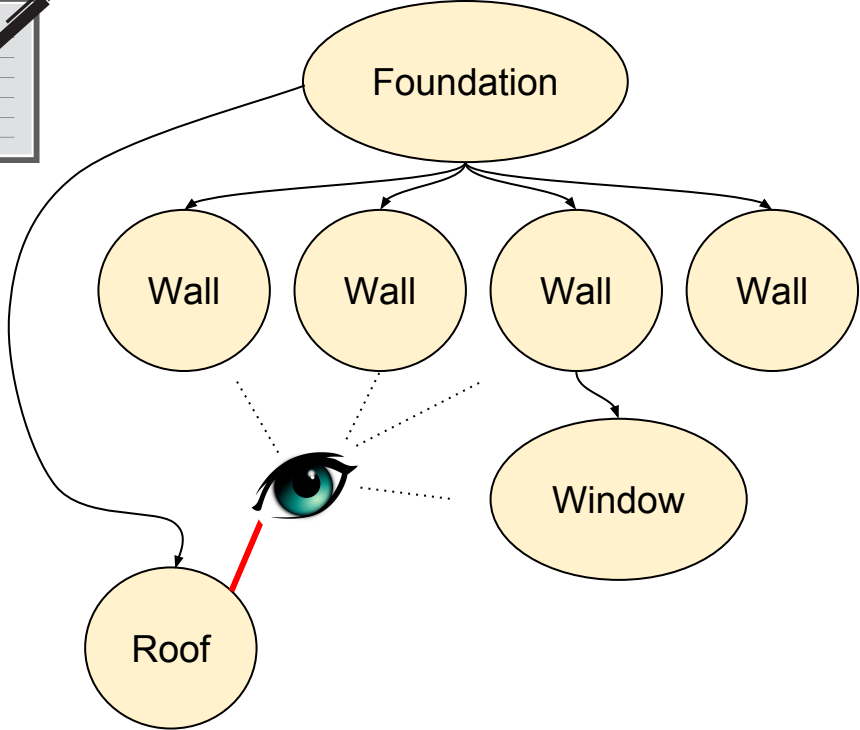
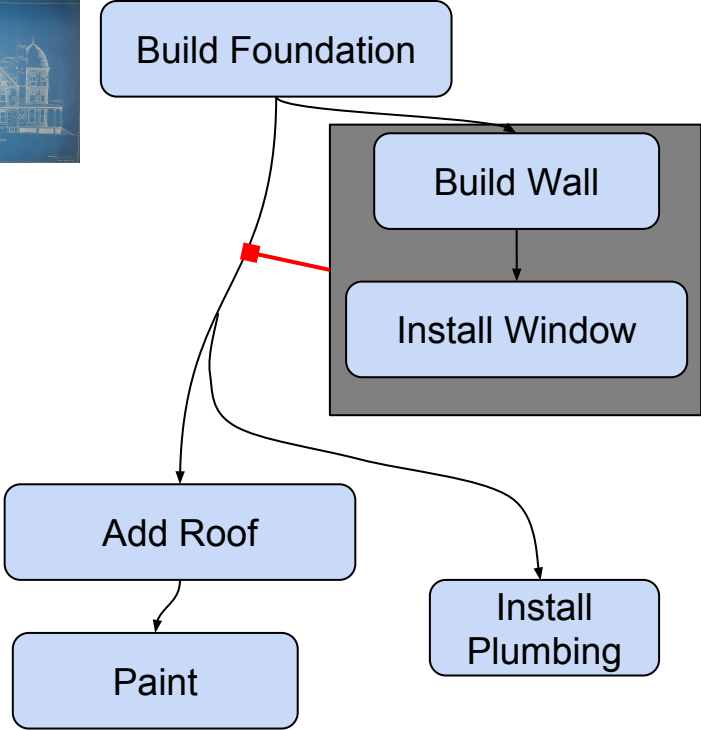
Workers have a time limit:

- Default 1 hour
- After reaching limit, workers exit
- However, a worker will never abandon a job in progress

Overtime makes the boss grumpy



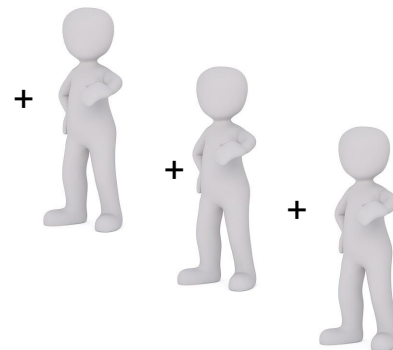
eHive fundamentals: walkthrough



eHive fundamentals: walkthrough



The role of the supervisor

The supervisor looks at the current job list, blueprints, and staff, and hires more workers if needed.



eHive fundamentals: walkthrough

A bit about specialization



- A worker knows about its capabilities.
- When looking for a job to claim, it checks the requirements for that class of job.
- If the worker's capabilities meet the job's requirements, it claims that job.
- A worker doesn't know "I can install lightbulbs." It knows "I can reach high" and it sees "install lightbulbs jobs have a "can reach high" requirement, so it knows it can specialize

eHive fundamentals: walkthrough

There has to be at least one job to get started

So there's a special mechanism to create at least one job when starting the project - usually this is the first class of job, but it can start anywhere

eHive fundamentals: components

How this model is actually implemented in eHive


eHive fundamentals: components



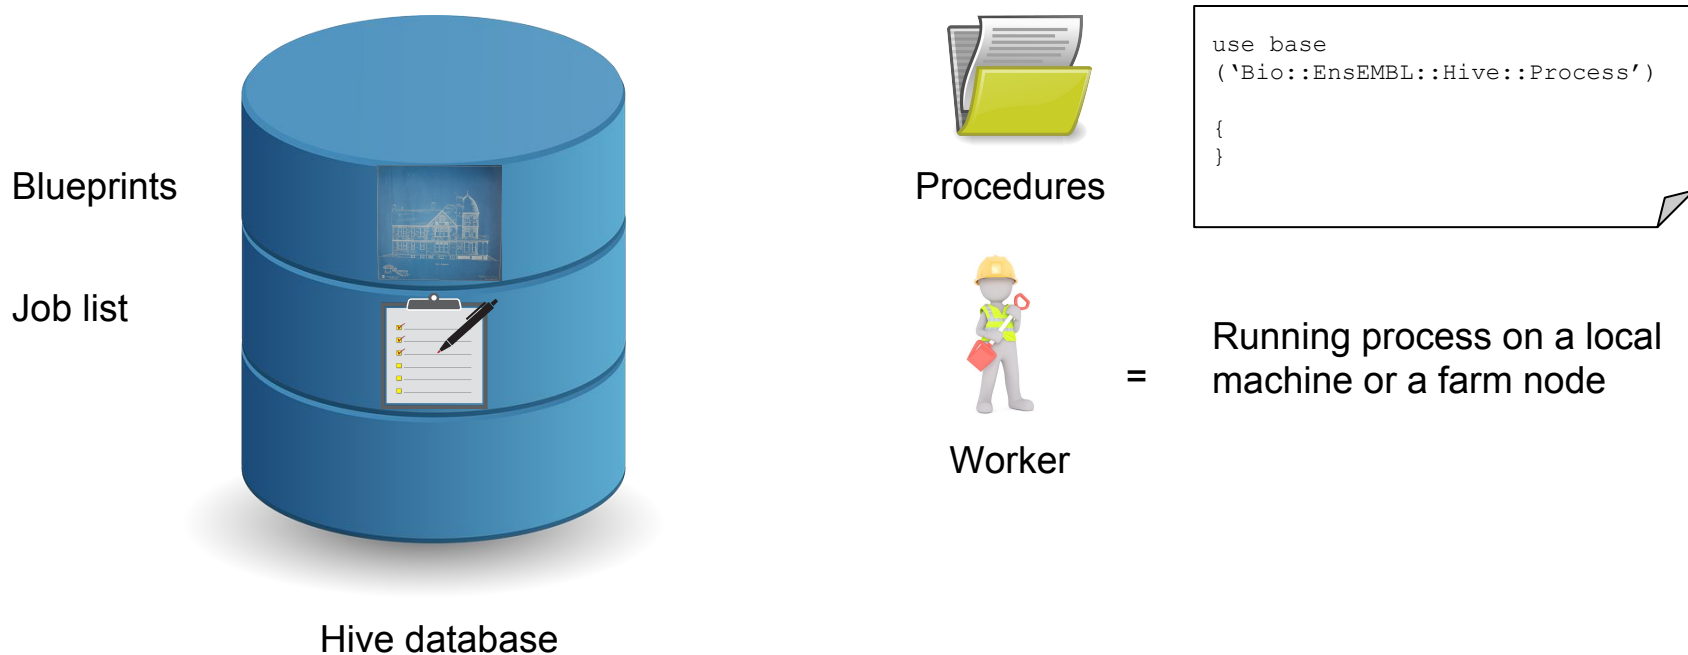
Can be run in any of:



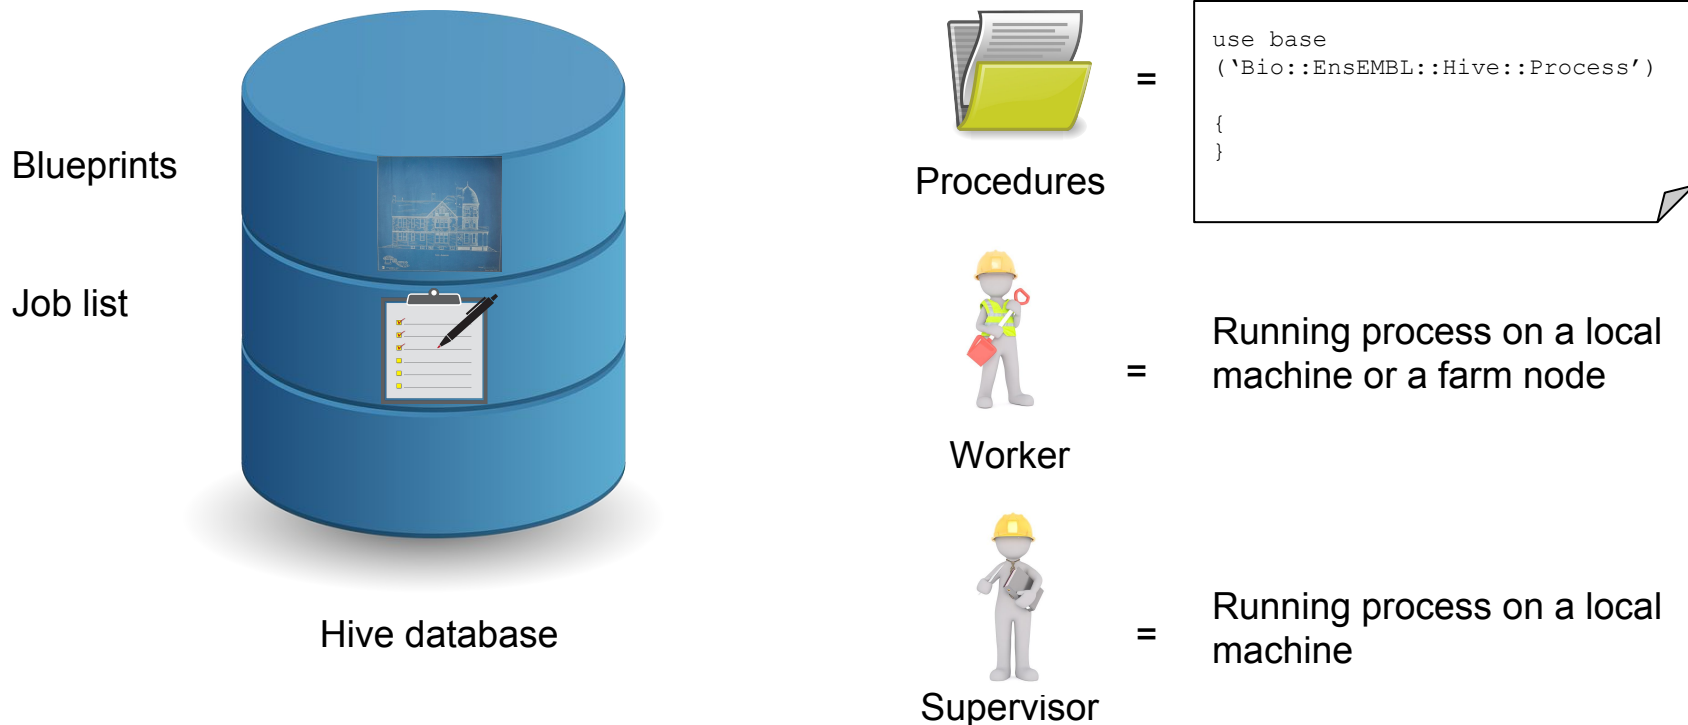
eHive fundamentals: components



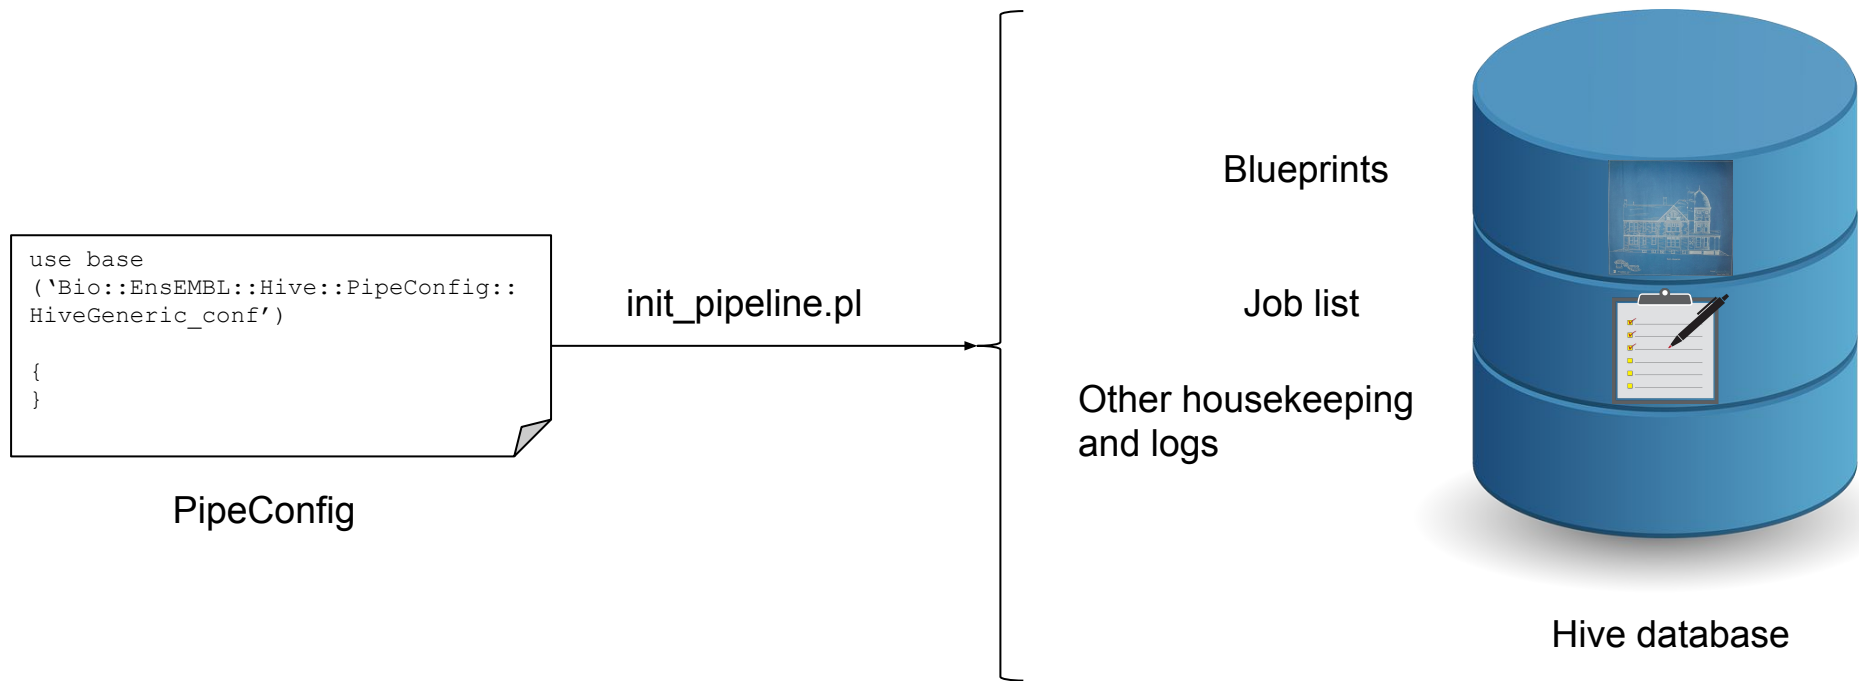
eHive fundamentals: components



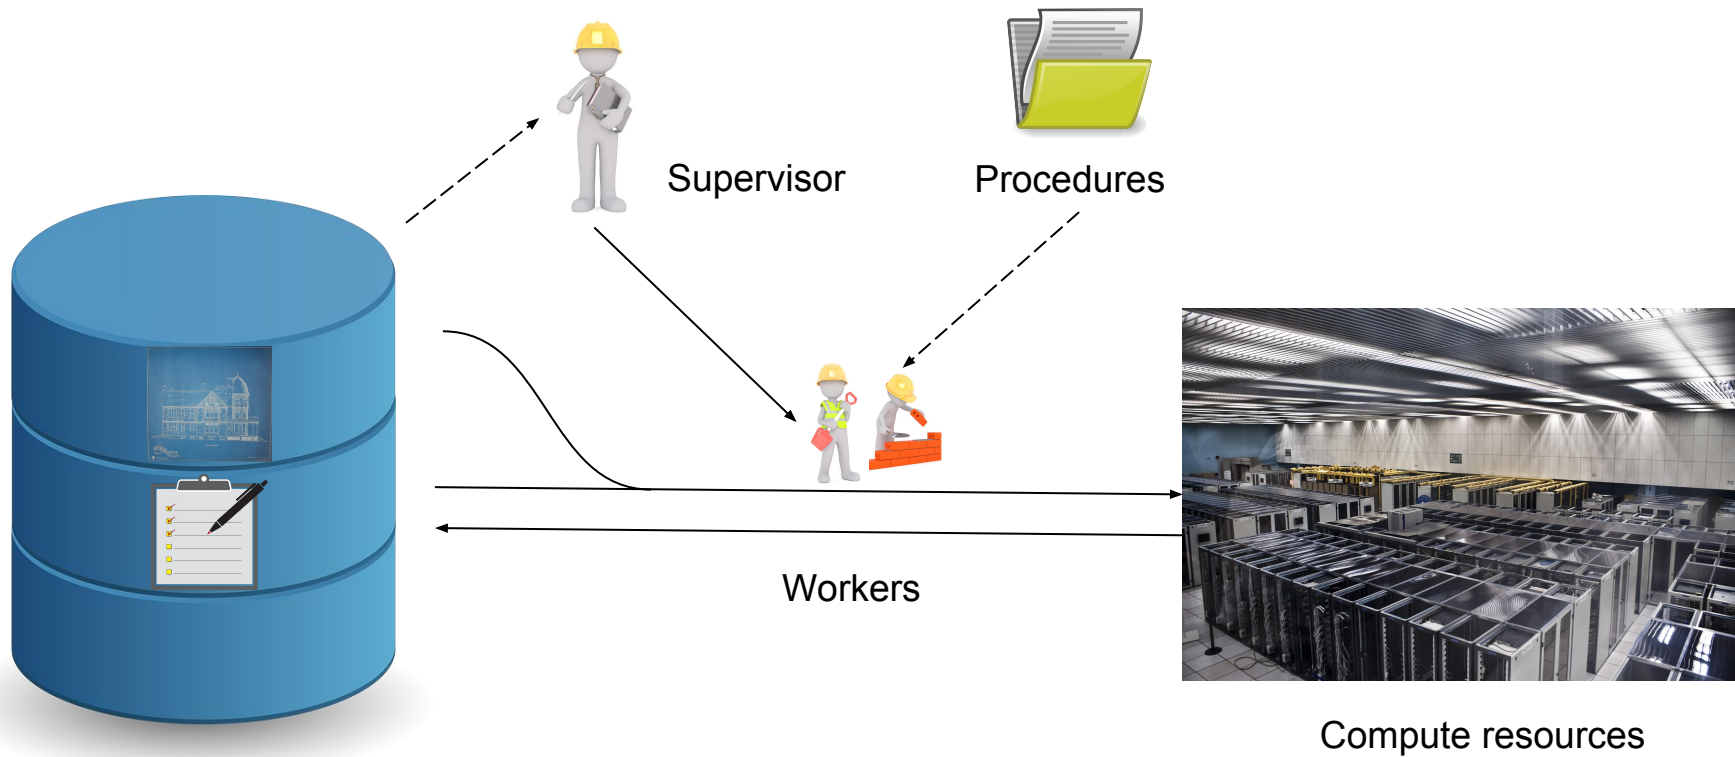
eHive fundamentals: components



eHive fundamentals: hive database



eHive fundamentals: running



eHive fundamentals: terminology

Blueprints stored
in a file



PipeConfig

Supervisor



Beekeeper

Worker



Worker

Procedure



Runnable

Compute resources



Meadow (~ a farm head node) and Valley (group of meadows)

Hive database



Hive database (details covered later in the course)

CC image by F. Hirzinger

Ensembl Acknowledgements

The Entire Ensembl Team

Bronwen L. Aken¹, Premanand Achuthan¹, Wasiu Akanni¹, M. Ridwan Amode¹,
Friederike Bernsdorff¹, Jyothish Bhai¹, Konstantinos Billis¹, Denise Carvalho-Silva¹,
Carla Cummins¹, Peter Clapham², Laurent Gil¹, Carlos García Girón¹, Leo Gordon¹,
Thibaut Hourlier¹, Sarah E. Hunt¹, Sophie H. Janacek¹, Thomas Juettemann¹,
Stephen Keenan¹, Matthew R. Laird¹, Ilias Lavidas¹, Thomas Maurel¹, William McLaren¹,
Benjamin Moore¹, Daniel N. Murphy¹, Rishi Nag¹, Victoria Newman¹, Michael Nuhn¹,
Chuang Kee Ong¹, Anne Parker¹, Mateus Patricio¹, Harpreet Singh Riat¹, Daniel Sheppard¹,
Helen Sparrow¹, Kieron Taylor¹, Anja Thormann¹, Alessandro Vullo¹, Brandon Walts¹,
Steven P. Wilder¹, Amonida Zadissa¹, Myrto Kostadima¹, Fergal J. Martin¹,
Matthieu Muffato¹, Emily Perry¹, Magali Ruffier¹, Daniel M. Staines¹, Stephen J. Trevanion¹,
Fiona Cunningham¹, Andrew Yates¹, Daniel R. Zerbino¹ and Paul Flicek^{1,2,*}

¹European Molecular Biology Laboratory, European Bioinformatics Institute, Wellcome Genome Campus, Hinxton, Cambridge CB10 1SD, UK and ²Wellcome Trust Sanger Institute, Wellcome Genome Campus, Hinxton, Cambridge, CB10 1SA, UK

Funding



Co-funded by the European Union