university of groningen / center for information technology

High Performance Computing

CIT

University of Groningen Site Talk

11th EasyBuild User Meeting
23-04-2026 | Guimarães, Portugal

Pedro Santos Neves

Short introduction

- HPC Team University of Groningen since 2023
- Background in Modeling in (Island) Evolutionary Biology
- Scientific Programmer in Evolutionary Biology @ University of Groningen 2022-2023
- EESSI / MultiXscale
- User support, training, software development

university of groningen / center for information technology

Projects @ CIT-RUG

- EESSI / MultiXscale
  - CernVM-FS Stratum 0

- Euclid

- Digital Lab

RUG HPC - Hábrók

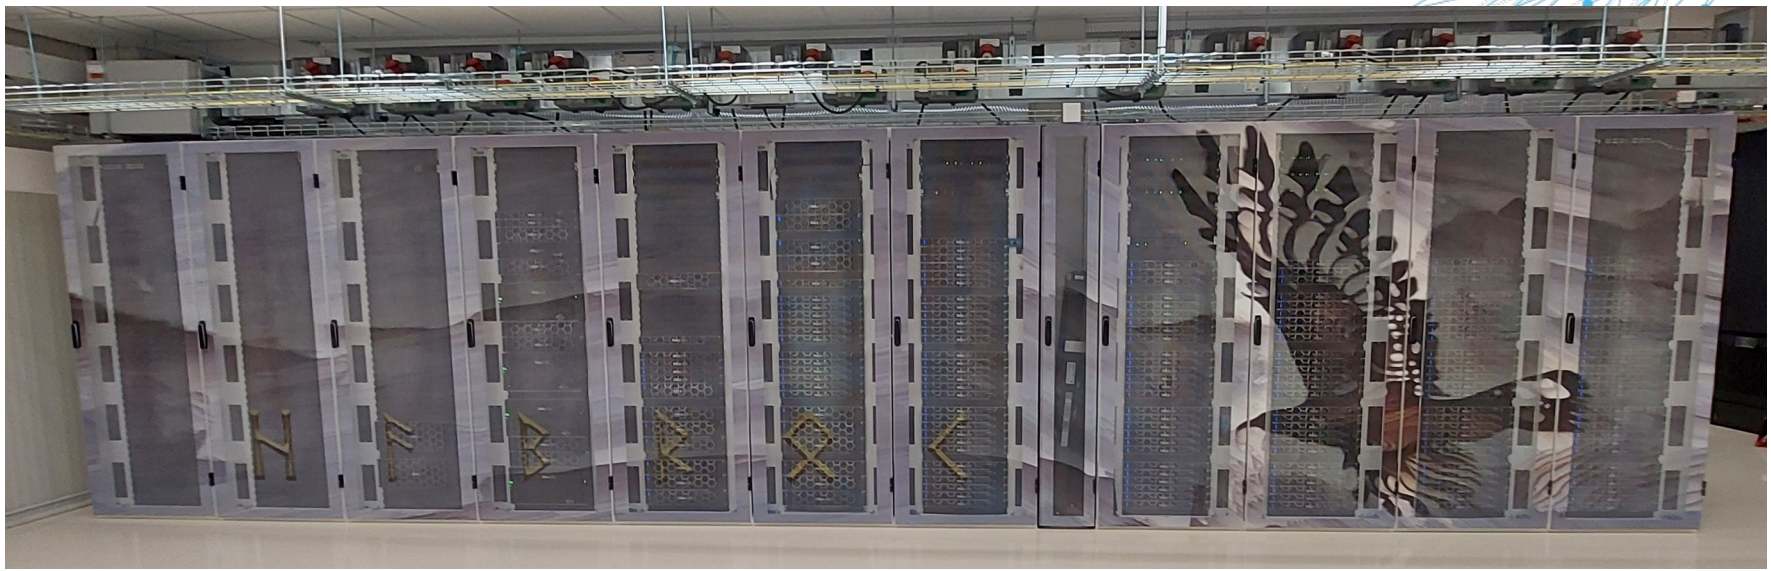

RUG HPC - Hábrók

university of groningen / center for information technology

Overview of Hábrók

- Operational since 2022
- Alma Linux 8 → 9
- 132 regular + parallel: AMD 7763 (Zen3) 512 GB memory
- 4 himem: Intel 8380 (Icelake) 4 TB memory
- 2 interactive GPU nodes: 1 x L40S, Intel 6240R (Cascadelake) 768 GB memory
- 3 GPU nodes: 8 x RTX Pro 6000, AMD EPYC 9575F (Zen5) 1.5 TB memory
- 6 GPU nodes: 4 x A100, Intel 8358 (Icelake) 512 GB memory
- 18 GPU nodes: 2 x V100, Intel 6150 (Skylake) 768 GB memory
- “restricted” nodes: 4 x H100, 2 x L40S owned by groups/projects
- 2.5 PB storage in Lustre filesystem

Software stack (1)

- Fully built with EasyBuild from the start
- Available through CernVM-FS
- Handful of licensed apps on NFS
- All modules exposed via default meta-module “2023.01”
- Built for Zen3, Icelake, Skylake
- Not optimizing builds for Cascadelake and Zen5 (until EESSI!)
- EESSI already available for a long time

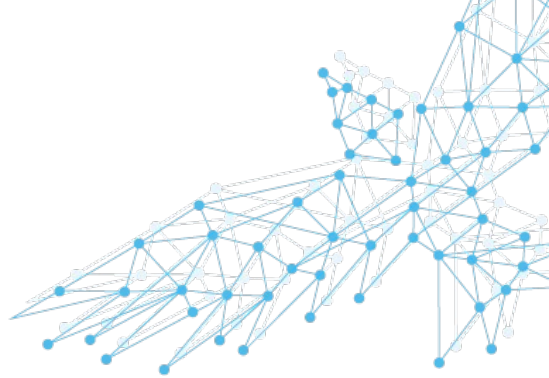

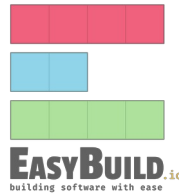

Software stack (2)

Built using private EESSI bot instance

- Uses custom build container based on Rocky 8 image
- Runs on service account in private GitHub repository
- Creates tarballs that can be ingested to Stratum 0
- Ingestion procedure is manual (for now)

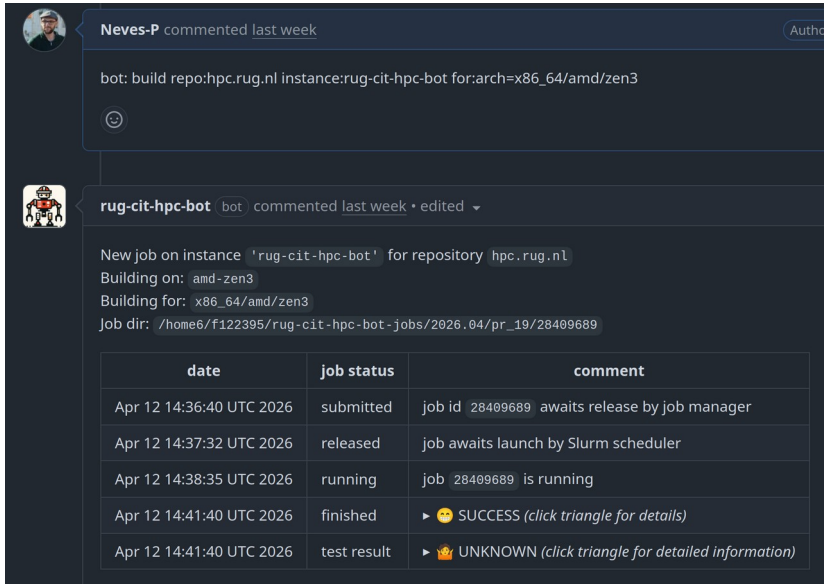

Software stack (3)

- Site specific code in local GitLab instance
  - Build container
  - Custom easyconfigs/easyblocks
  - All software tracked on easystack file
- Moving the bot to GitLab would make things even better

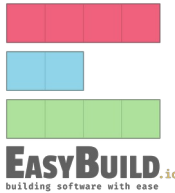

Challenges & surprises

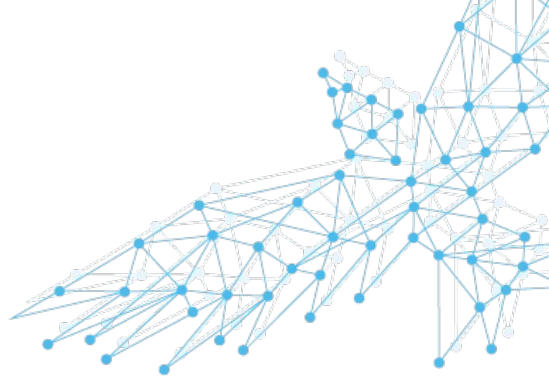

- Quite heterogeneous cluster
  - Oldest nodes purchased in 2019 (previous cluster) V100
  - Newest nodes commissioned in 2026: RTX PRO 6000
  - Some hardware out of support (V100 – drivers, compute capabilities)
  - Without EasyBuild, managing software stack would have been so much harder
- Not providing optimized builds for RTX PRO 6000 nodes (yet); EESSI will help

Challenges & surprises - Whisper

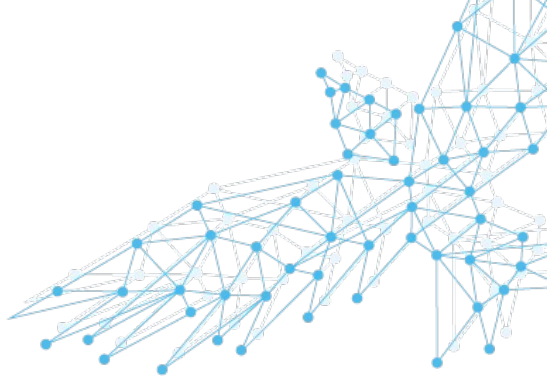

- Popular open source transcription and translation tool by OpenAI

- >600 users in ~6 months (on top of usual rate of new users)
- Deployed as custom Open OnDemand App (Job Composer)
- Defaults to V100 nodes
- Low barrier of entry for non-technical users
- Not a good fit for traditional HPC systems / batch scheduling
  - Dedicated infrastructure needed!

Challenges & surprises - Whisper

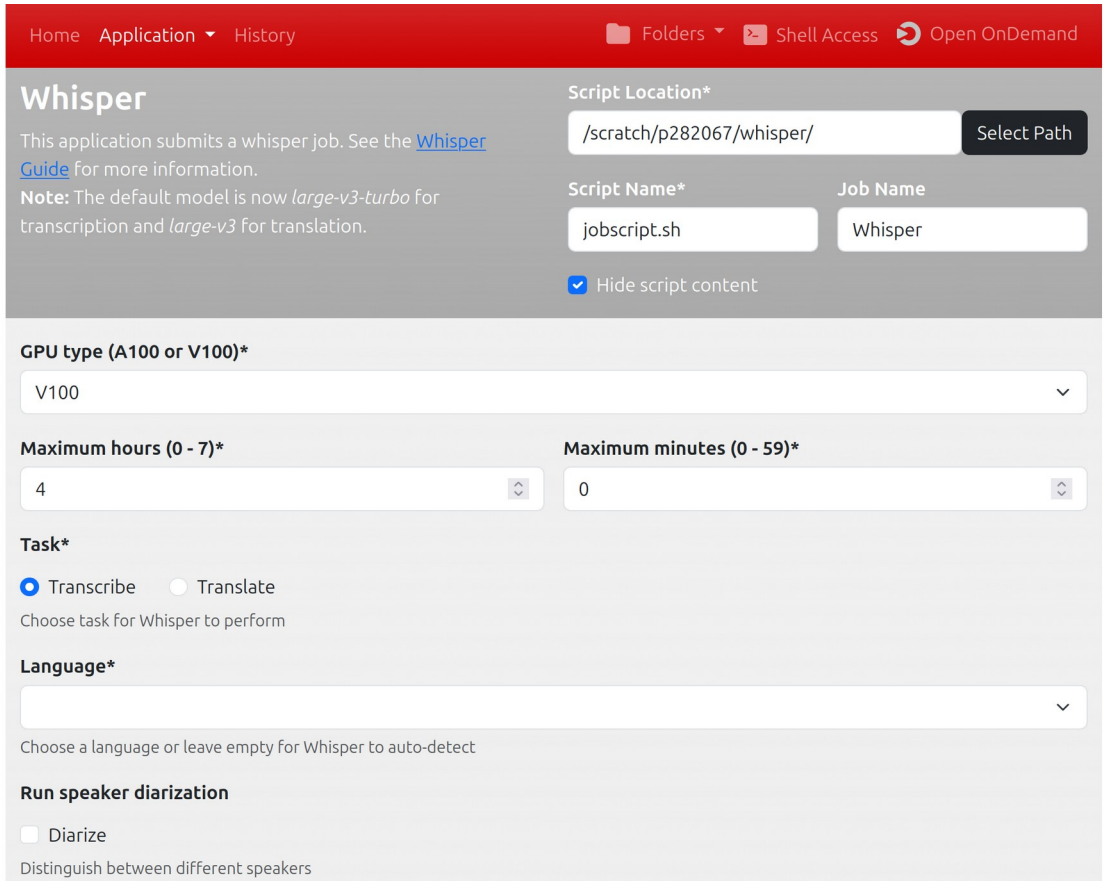

university of groningen / center for information technology

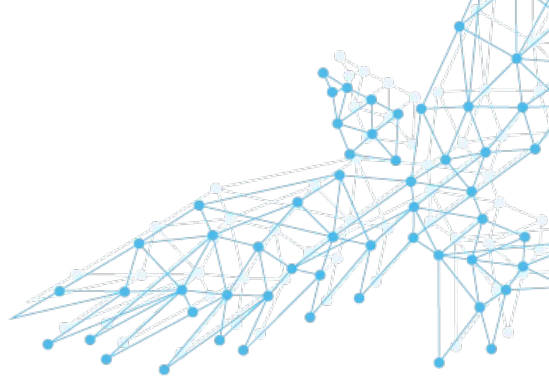

Next steps

Next steps – Phased change to Alma 9

- Deploy subset of nodes with Alma 9
  - Testing indicates most (all?) software stack works
  - Interactive nodes
- Allocate Alma 9 nodes to reservation and let users test
- After initial phase, upgrade some interactive nodes to Alma 9
- New software **only built on top of EESSI**

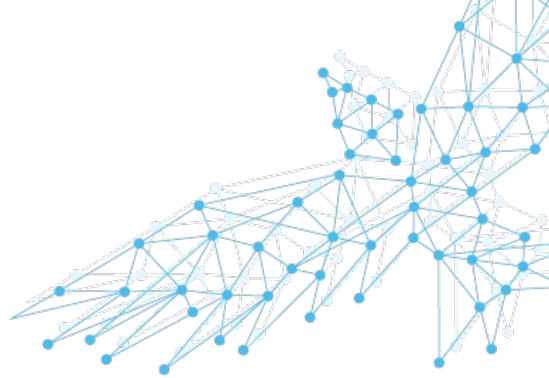

EESSI
EUROPEAN ENVIRONMENT FOR SCIENTIFIC SOFTWARE INSTALLATIONS

Alma Linux upgrade & adopting EESSI

- Before
  - Alma Linux 8
  - Custom software stack + side mounted EESSI
- 2940 total modules per CPU target
  - We could never have managed without EasyBuild!

- After
  - Alma Linux 9
  - EESSI as base stack + `EESSI-extend` built modules
- Much reduced number of custom modules
- Old software stack still available for compatibility
- Now software only via EESSI / `EESSI-extend`

Recent and ongoing work – EESSI as base stack

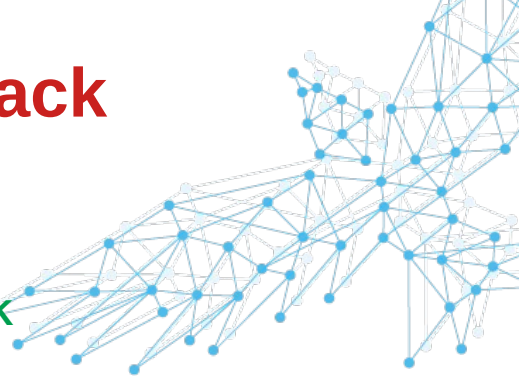

- EESSI has been available at RUG for a long time
  - `setenv EESSI_MODULE_FAMILY_NAME software_stack`
- Write documentation on using EESSI stack
  - Rework training materials
  - Mitigate and catch issues at source (e.g., users who run self-compiled code)
- `EESSI-extend` configuration
- Major milestone:
  - Rework CernVM-FS server deployment

CernVM-FS setup - before

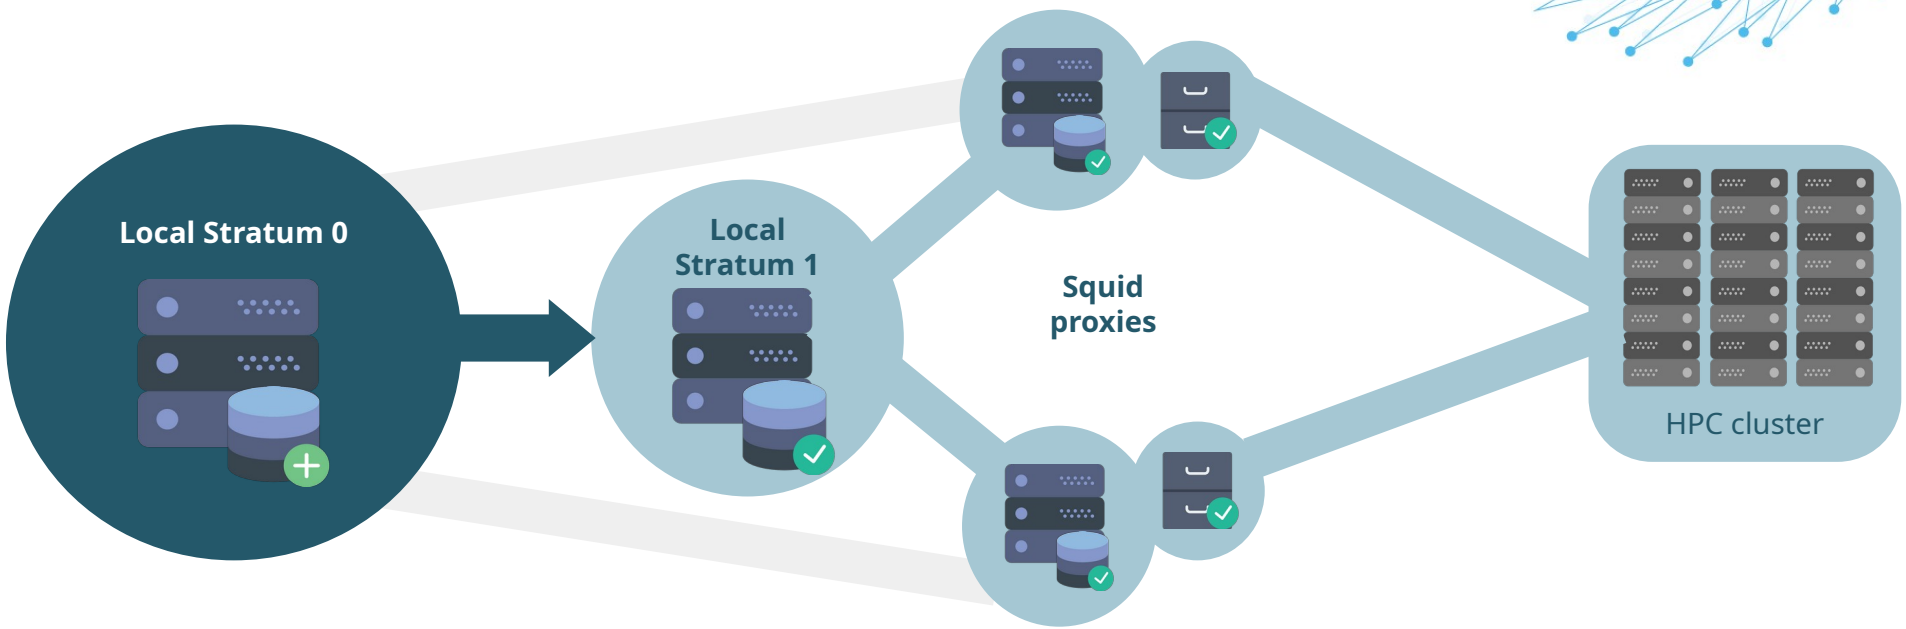

2 cores
4 GB RAM
500 GB disk

3 cores
16 GB RAM
500 GB disk

3 cores
16 GB RAM
100 GB disk

university of groningen / center for information technology

CernVM-FS setup - now

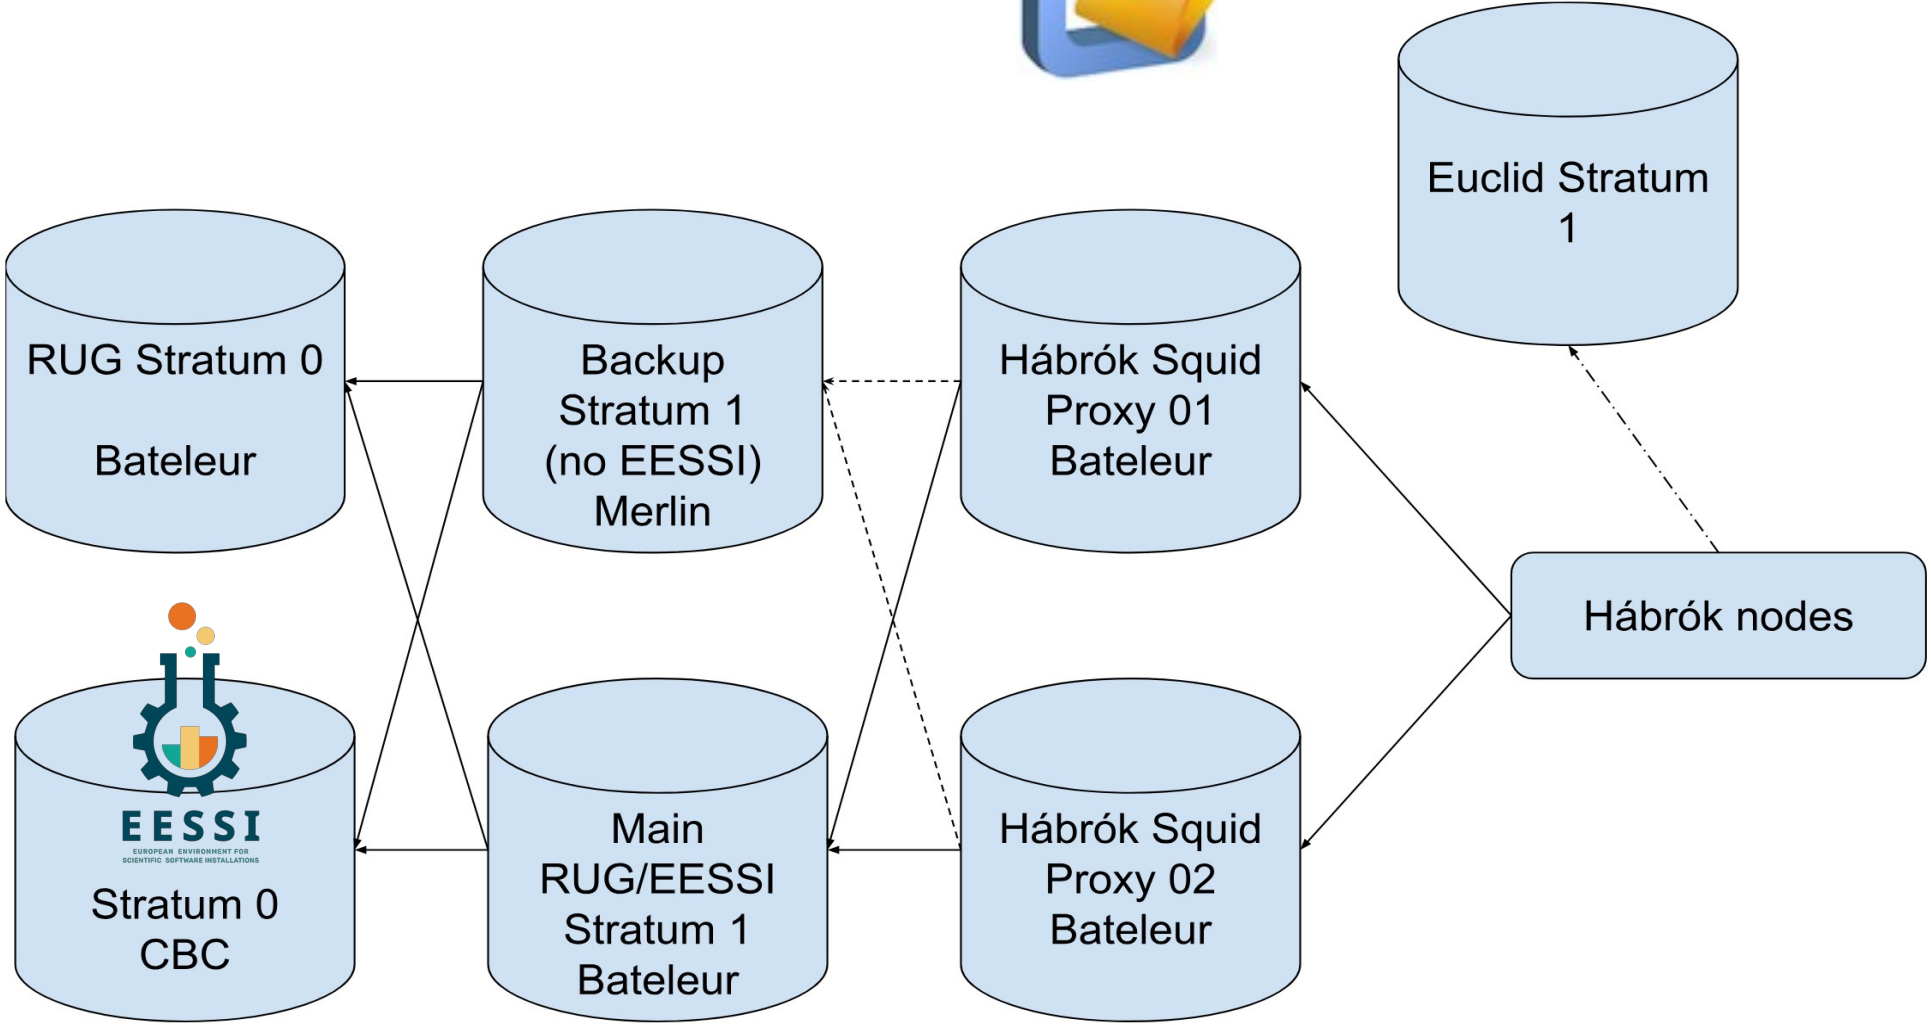

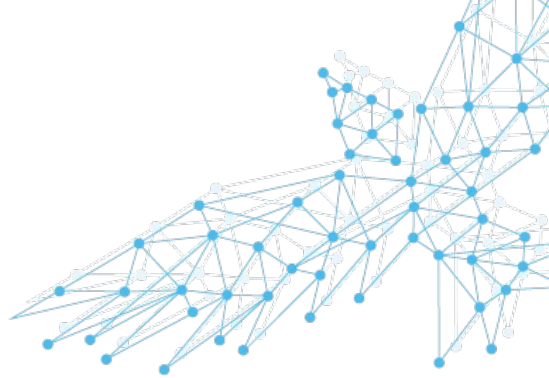

Thank you!
Questions?

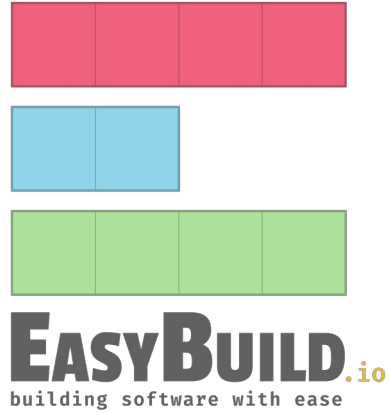

university of groningen / center for information technology