

MUNI

New technology course at FI MU

Tomáš Rebok and Lukáš Hejtmánek
Institute of Computer Science
Masaryk University



Course idea

Why a new course?

- to fill the gap between
 - theoretical concepts and common computing approaches
 - e-infrastructures and HPC
- **Question to the audience:** what do you expect from such a course?
 - which technologies, concepts, etc.
 - let's have a brainstorm...

Course idea II.

What we want to deliver?

- overview of various e-infrastructure technologies
 - including their fundamental concepts
- overview of available national e-infrastructure technologies and resources
 - including their practical usage
- application of selected problems from different research areas to e-infrastructures (across the presented technologies)
 - artificial intelligence, data analysis, etc.
 - may change in time

Course idea III.

Basic course setup

- theoretical lectures – 2 study hours every week
- practical trainings – 1-2 hours every week (via webinar)
 - bi-weekly – training + consultations

Types of completion

- examination – written test + project
- colloquium – project
- zápočet – Bc students

Course overview

Course content (1st version):

- Introduction to HPC infrastructures in e-INFRA CZ
 - CERIT-SC
 - MetaCentrum
 - IT4I
- Introduction to HPC technologies
 - Kubernetes and Container technologies
 - OpenStack Cloud
 - OpenPBS (Slurm?)
- Storage (CESNET)
 - Object storage, Ceph, S3
- Efficient computing
 - GPU computing
 - Pilot jobs – mass workload processing
 - Data-processing pipelines
 - OneData – unified data management and access

Infrastructures

- e-INFRA CZ infrastructures overview and missions
- Supported technologies in e-INFRA CZ
 - CERIT-SC, MetaCentrum, IT4I
- Resources available
- Access
- System of user accounts
- User support
- European e-infrastructure space (EGI.eu)

Kubernetes and Container technologies

- About 3 lectures
 - Docker, docker-compose, building docker images, docker registry
 - Kubernetes principles, architecture overview
 - Kubernetes networking, various add-ons (certificates), application deployment, helm packages
- Students should learn
 - How to build a docker image, how to run it
 - How to deploy a small application
 - How to debug the deployment

OpenStack Cloud

- 3 lectures on IaaS type of cloud service
 - Virtualization, VM images, configuration of VMs in the cloud
 - and using OpenStack to build virtual infrastructures, deploy services
 - Principles and architecture of clouds built on OpenStack
 - overview of important OS (micro)services, how it is connected
 - HW deployment
 - Strategies of building private cloud service (with example beskar.cloud)
- Students should learn
 - Deploy and configure a service on virtual machines using EC2 compatible cloud service and various orchestrating tools – API/CLI, Terraform
 - Principles of virtualization and benefits of deploying/operating private-cloud
 - Simple debugging of deployed VMs

OpenPBS

- ca 2 lectures
- OpenPBS basics
 - the idea of batch computing, jobs (interactive, batch)
 - includes some historical view about how these concepts evolved
 - infrastructure resources and association with jobs
 - jobs – submissions and monitoring
 - data handling – scratches (local, shared), stagein & stageout
- OpenPBS advanced
 - parallel and distributed computing (OpenMP & MPI)
 - special use-cases – job dependencies, checkpointing, job arrays, etc.
 - PBS & containers (Singularity)

Storage

- About 1 lecture
 - How to efficiently use storages connected to e-INFRA CZ
 - including the support of teamwork and securing of the data
 - NFS servers
 - How to move the data efficiently (and securely)
 - to/from/inside/outside the infrastructure
 - Cesnet object storage
 - CEPH
 - S3
 - Working with sensitive data (SensitiveCloud)

Efficient computing

- 2 lectures
- GPU Computing
 - efficient use of GPUs in the e-infrastructure
 - TBD with Fila 😊
- Pivot Jobs – mass workload processing (dynamic task scheduling)
 - efficient processing of large number of datafiles in e-infrastructure
- Data-processing pipelines
 - workflow-based processing
 - workflow automation (orchestrators – AirFlow, Argo, Galaxy)
 - including their use-cases (AFoLab)
- OneData – unified data management and access

System of Labs

- Bi-weekly in the form of a Webinar – about 6 labs
 - Consultation hours
- Topics
 - Docker
 - Kubernetes 1
 - Kubernetes 2
 - OpenStack
 - PBS
 - Storage

Docker and Kubernetes

- Docker
 - Creating docker image and docker-compose script
- Kubernetes 1
 - Simple deployment of the image above, using docker registry
- Kubernetes 2
 - Advanced deployment, running jobs from the deployment
 - Helm package

OpenStack

- Lab: VM orchestration using Terraform and configuration with Ansible
 - Running simple web server application
 - Networking – publishing app to the Internet, managing ingress/egress policies
 - Simple load balancing, high-availability service
 - Using software defined storage for block devices or S3

PBS

- Running simple and advanced jobs
- Using scratch
- Distributed computing, array of jobs
- Using singularity for running container

Storage

- Deploying minio server (OpenStack, Kubernetes)
- Upload/download data
- Setting access policy

Lab Evaluation

- Automatic tests
 - Binary evaluation pass/fail
- Plagiarism?

Final Exam

- Written test
- Example questions:
 - CloudEdge Computing Inc. hired you to develop cloud-enabled mobile applications. These applications have strict timing requirements that require cloud-edge capabilities beyond what mobile devices alone can offer. For users that are constantly on the move and expect you to maintain high quality of service, your application running on the edge will need to frequently migrate from one edge to another. Would you recommend container or VM-based deployments on the edge? Justify your recommendation. (Points: 4)
 - What should the mascot be for this class? Why? (Points: 1)
 - Answer: Chameleon – 2 points.

MUNI

Thank you.

