



Container platform Kubernetes

Lukáš Hejtmánek


CERIT-SC



Kubernetes Platform Services

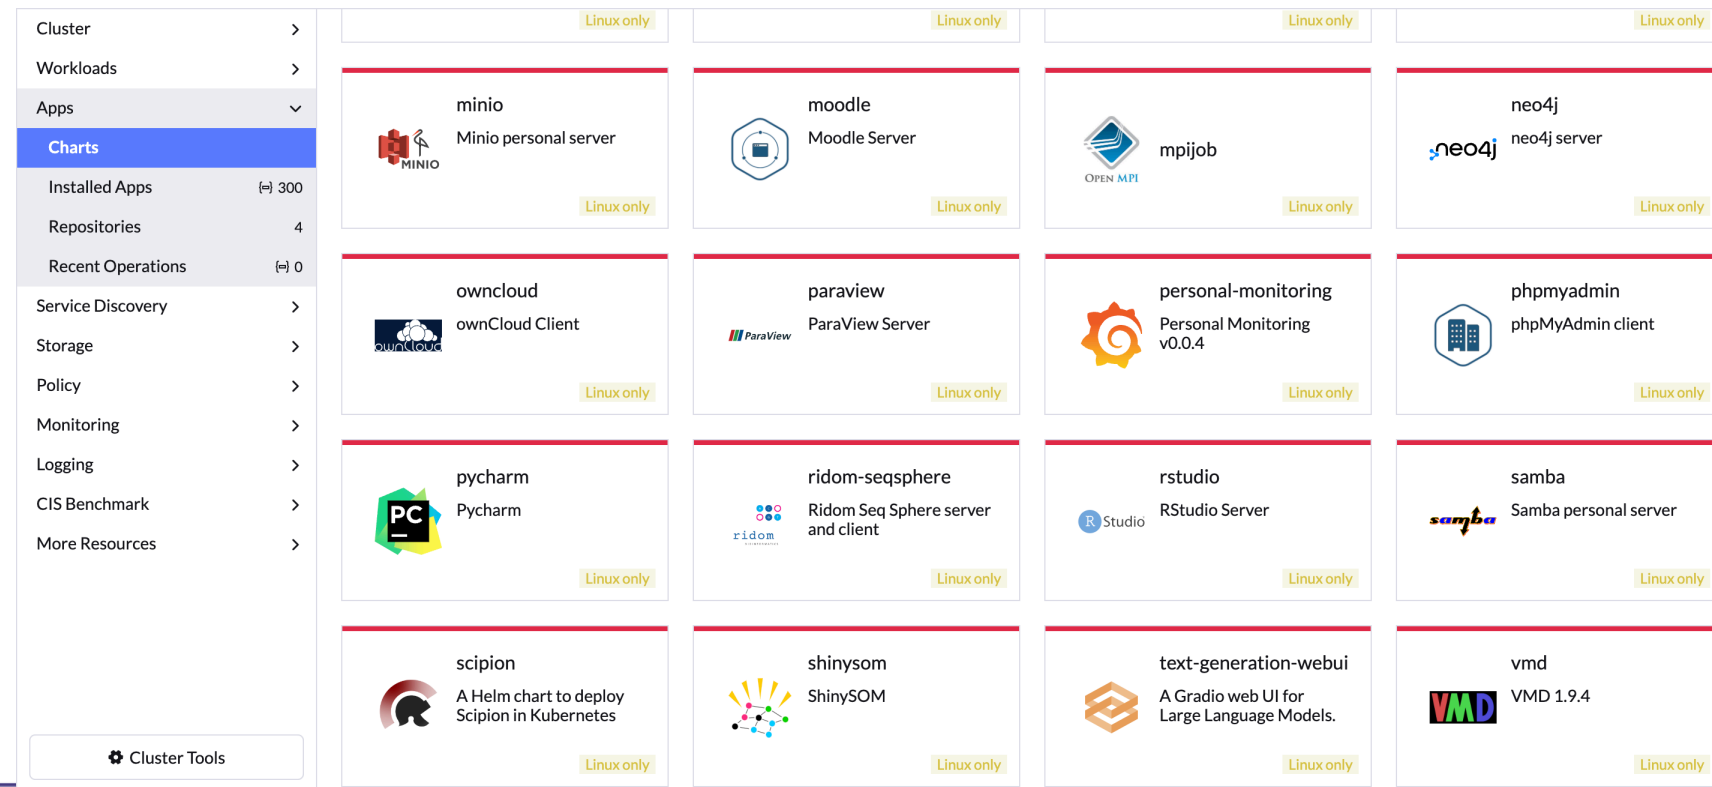


- Rancher application catalogue
- Web applications
- Native Kubernetes API



Rancher Applications

- Catalogue of applications — mostly remote GUI
 - Ansys, Matlab, Rstudio, Generic Desktop, VMD
 - Code-Server, Pycharm
 - UI for LLM

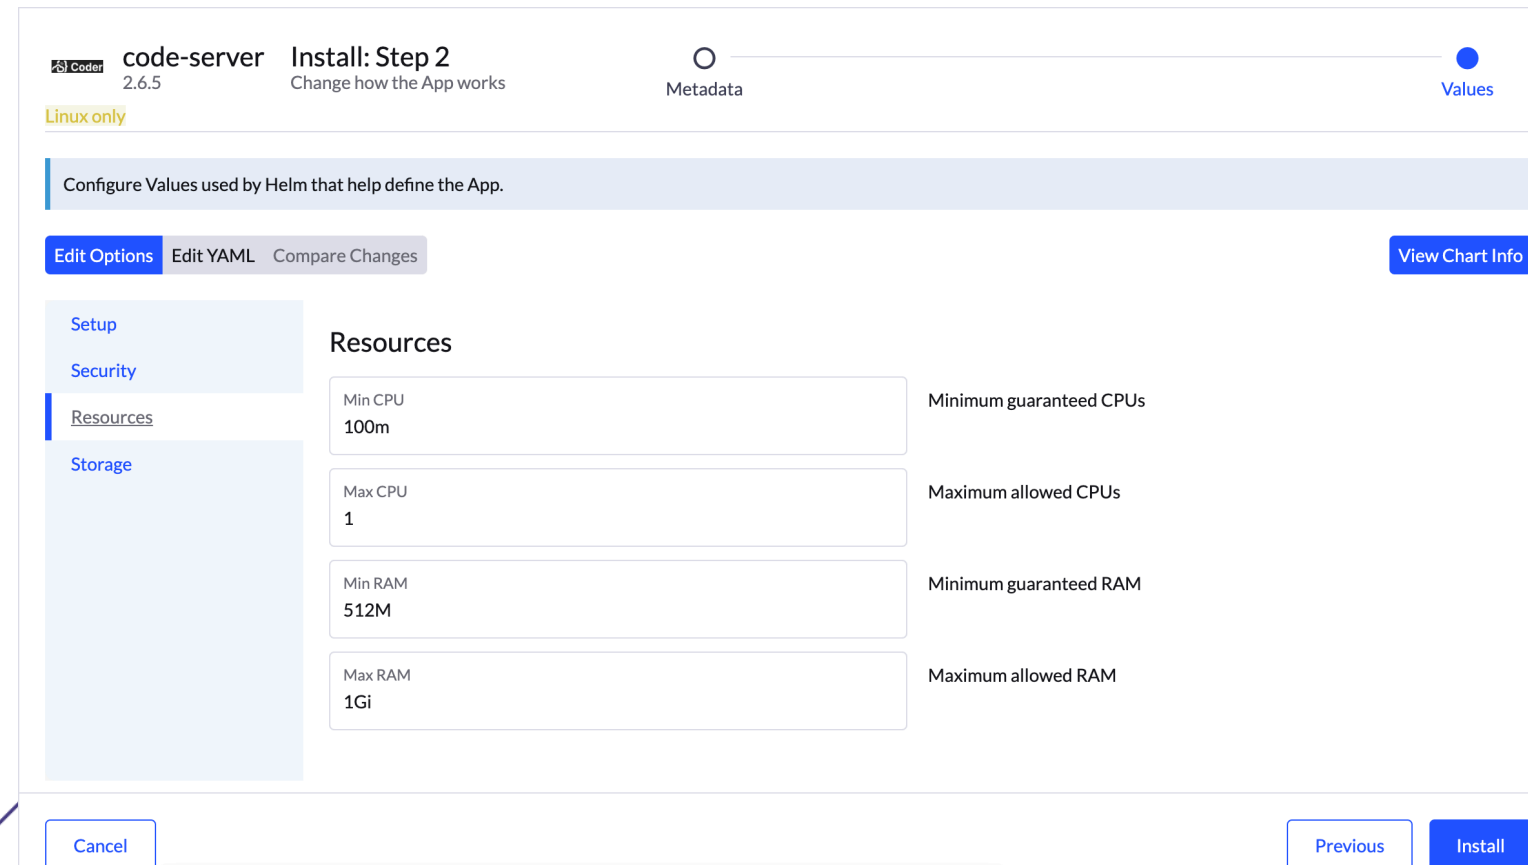


The screenshot displays the Rancher UI interface. On the left is a navigation sidebar with the following items: Cluster, Workloads, Apps, Charts (highlighted in blue), Installed Apps (300), Repositories (4), Recent Operations (0), Service Discovery, Storage, Policy, Monitoring, Logging, CIS Benchmark, and More Resources. At the bottom of the sidebar is a 'Cluster Tools' button. The main area shows a grid of application cards, each with an icon, name, description, and a 'Linux only' label. The cards are:

Application Name	Description	Platform
minio	Minio personal server	Linux only
moodle	Moodle Server	Linux only
mpijob	Open MPI	Linux only
neo4j	neo4j server	Linux only
owncloud	ownCloud Client	Linux only
paraview	ParaView Server	Linux only
personal-monitoring	Personal Monitoring v0.0.4	Linux only
phpmyadmin	phpMyAdmin client	Linux only
pycharm	Pycharm	Linux only
ridom-seqsphere	Ridom Seq Sphere server and client	Linux only
rstudio	RStudio Server	Linux only
samba	Samba personal server	Linux only
scipion	A Helm chart to deploy Scipion in Kubernetes	Linux only
shinyson	ShinySOM	Linux only
text-generation-webui	A Gradio web UI for Large Language Models.	Linux only
vmd	VMD 1.9.4	Linux only

Rancher Applications — How to Use?

- Log to rancher: <https://rancher.cloud.e-infra.cz>
- Select cluster, Apps, Charts
- Fill option form (go through all options)
- Application should be running
- Navigate to Service
- See <https://docs.cerit.io>



code-server 2.6.5 Install: Step 2
Change how the App works

Linux only

Configure Values used by Helm that help define the App.

Edit Options Edit YAML Compare Changes View Chart Info

Setup
Security
Resources
Storage

Resources


Min CPU 100m	Minimum guaranteed CPUs
Max CPU 1	Maximum allowed CPUs
Min RAM 512M	Minimum guaranteed RAM
Max RAM 1Gi	Maximum allowed RAM

Cancel Previous Install

Rancher Applications — Missing One?



I want **YOU** to request!
Requesting station:
k8s@ics.muni.cz






Web Applications

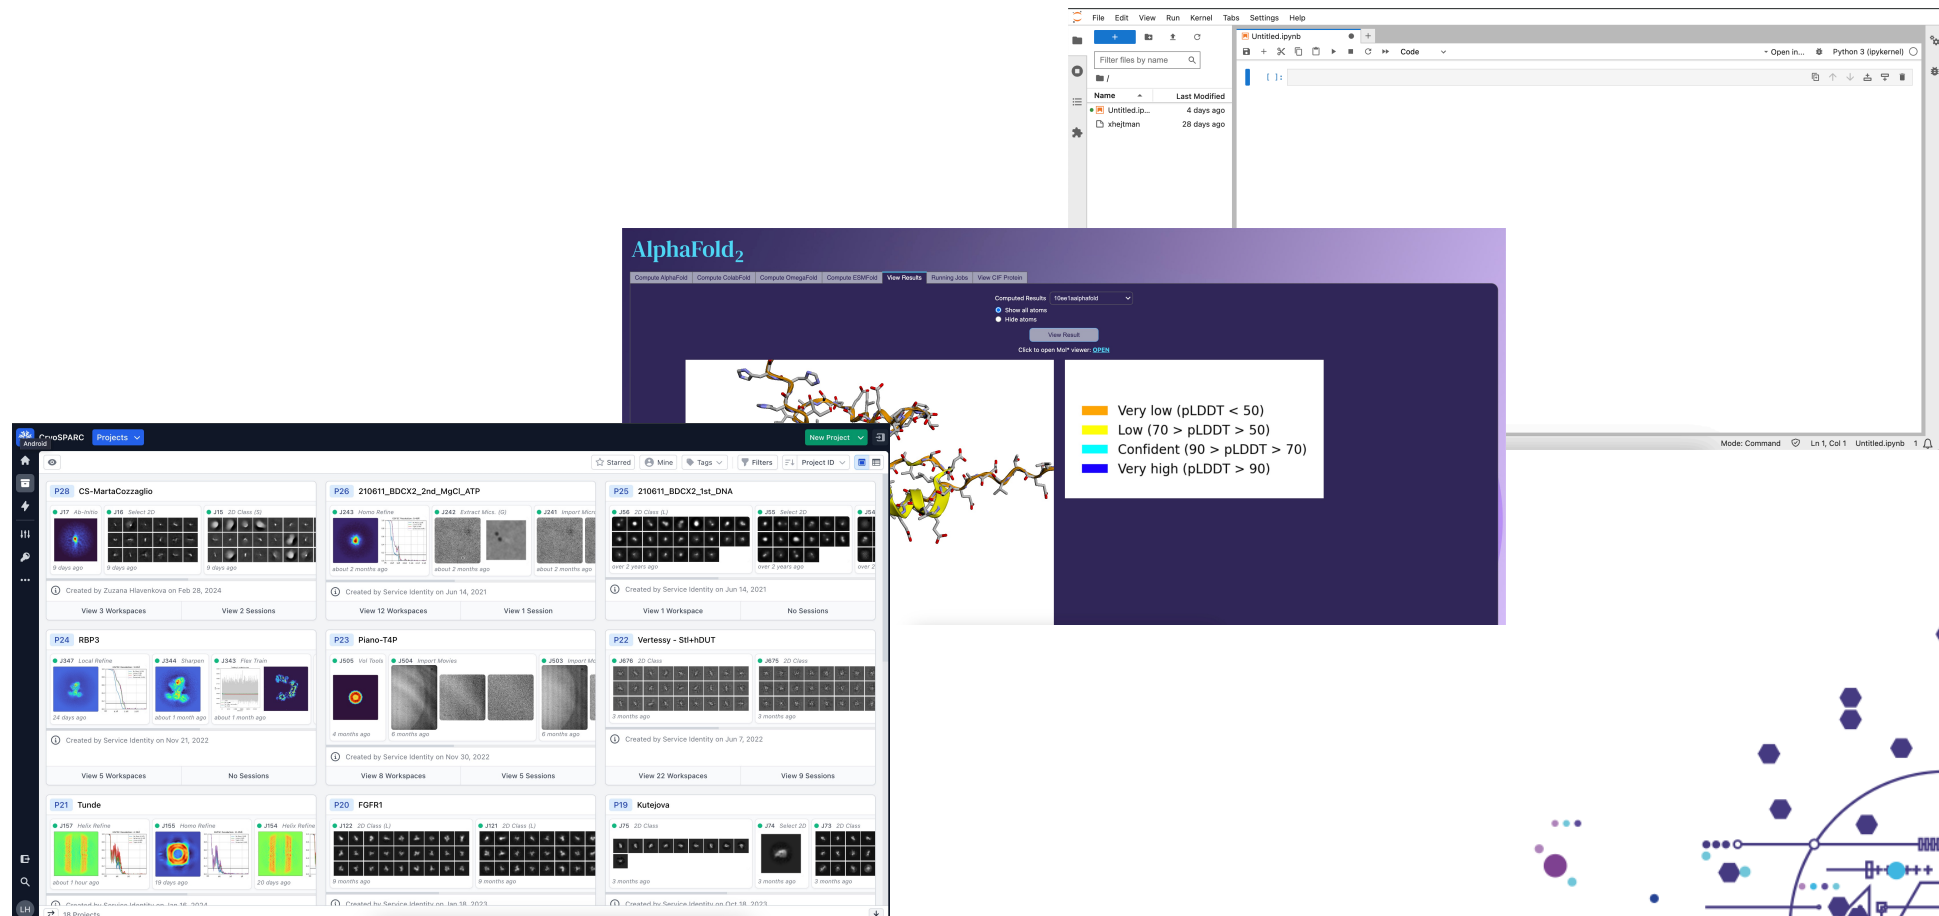


+



Web Applications

- Jupyterhub
- AlphaFold
- Cryosparc




The image displays a JupyterLab interface with a grid of project workspaces. Each workspace contains various data visualizations, including heatmaps and 3D molecular models. A prominent 3D molecular structure is shown in the center, with a legend indicating confidence levels for pLDDT values:

- Very low (pLDDT < 50)
- Low (70 > pLDDT > 50)
- Confident (90 > pLDDT > 70)
- Very high (pLDDT > 90)

The interface also shows a file browser on the left and a code editor on the right, indicating a full-featured development environment.

Jupyterhub

- <https://hub.cloud.e-infra.cz>
- Integrated with e-INFRA CZ storages
 - Integration with Onedata datasets in progress
- Without explicit run-time limit
- Number of prepared images — ask if you need more k8s@ics.muni.cz
 - Rstudio, Matlab, RSAT images
- Allows to spawn additional jobs from the notebook
- Notebooks for advanced users
 - <https://binderhub.cloud.e-infra.cz>
 - Builds notebook image from source
 - Possible to request resources
 - CPU, GPU, Memory



Server Options

Choosing image

Predefined Jupyter Notebook images are:

- Minimal NB - Jupyter minimal notebook
- Datascience NB - Jupyter datascience NB (Python, R, Julia)
- Scipy NB - Jupyter scipy NB (minimal NB + interactive visualizations and plots in Python)
- Tensorflow NB - Jupyter TensorFlow NB (scipy NB + tensorflow and keras)
- Tensorflow GPU NB - Jupyter TensorFlow NB on GPU (Jupyter Notebooks with the power of GPU, perform GPU calculations using Tensorflow and Pytorch in collaborative notebooks. Displayed device is XLA_GPU)
- RationAI with GPU, TF, TB - Jupyter TensorFlow NB on GPU (Image for specific team)
- RStudio - RStudio via Jupyterhub
- Alphapose - Alphapose AI framework

Apart from Jupyter Notebook images, you can also spawn RStudio instance with *R 4.2.1*. GPU is not supported in the image, all other features are support (including conda). If you seek another version of R with RStudio, GPU support or have other requirements send your inquiry to our [IT Service desk](#).

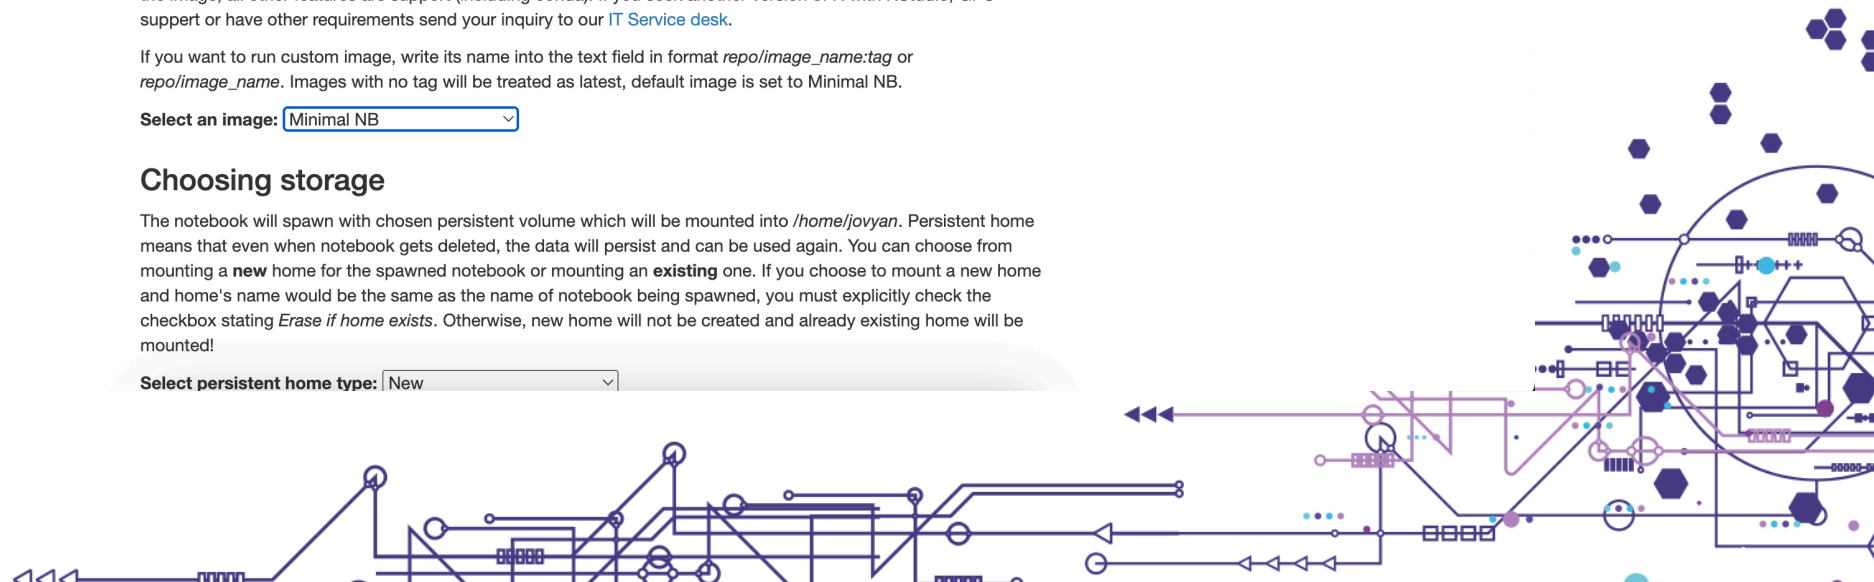
If you want to run custom image, write its name into the text field in format `repo/image_name:tag` or `repo/image_name`. Images with no tag will be treated as latest, default image is set to Minimal NB.

Select an image:


Choosing storage

The notebook will spawn with chosen persistent volume which will be mounted into `/home/jovyan`. Persistent home means that even when notebook gets deleted, the data will persist and can be used again. You can choose from mounting a **new** home for the spawned notebook or mounting an **existing** one. If you choose to mount a new home and home's name would be the same as the name of notebook being spawned, you must explicitly check the checkbox stating *Erase if home exists*. Otherwise, new home will not be created and already existing home will be mounted!

Select persistent home type:



- You can request to deploy a custom instance
 - Access limited to a Perun group
 - Custom start form and images



Server Options

Image

Select an image

Home

Erase if home exists

Consider thoroughly checking this option - it removes whole home directory and all data located there (/home/jovyan/). Use only when notebook is broken so it does not start. In other cases, remove data from terminal inside the notebook.

Resources

CPU

Please choose amount of CPUs which will be assigned to notebook (default 1):

Select number of CPU (1-32):


Memory

Please choose upper memory limit (in GB) which will be assigned to notebook (default 4):

Start

Alphafold

- Web UI for Alphafold tools
- Supports several different tools
 - Alphafold, Colabfold, Omegaold, Esmfold
- Colabfold
 - Running our mmseqs server (no rate limit)
 - Added possibility to limit number of models
- Integrated Mol* viewer and e-INFRA CZ storage (brno12)



AlphaFold₂

Compute AlphaFold | Compute ColabFold | Compute OmegaFold | Compute ESMFold | View Results | Running Jobs | View CIF Protein

Protein name	<input type="text" value="adventurous-frank"/>
Proteins	<input type="text" value=">Sequence1"/>
Max template date	<input type="text" value="default"/>
DB Preset	<input type="text" value="full_dbs"/>
Model Preset	<input type="text" value="monomer"/>
Precomp MSAS	<input type="checkbox"/>
Predictions per model	<input type="text" value="5"/>
Run Relax	<input checked="" type="checkbox"/>
Make results public	<input checked="" type="checkbox"/>
E-mail	<input type="text" value="xheitman@ics.muni.cz"/>
Force computation	<input type="checkbox"/>
Version of AlphaFold	<input type="text" value="Alphafold 2.3.1"/>

✓ Run

AlphaFold₂


Compute AlphaFold Compute ColabFold Compute OmegaFold Compute ESMFold **View Results** Running Jobs View CIF Protein

Computed Results 10ee1aalpafold

- Show all atoms
- Hide atoms

View Result

Click to open Mol* viewer: [OPEN](#)

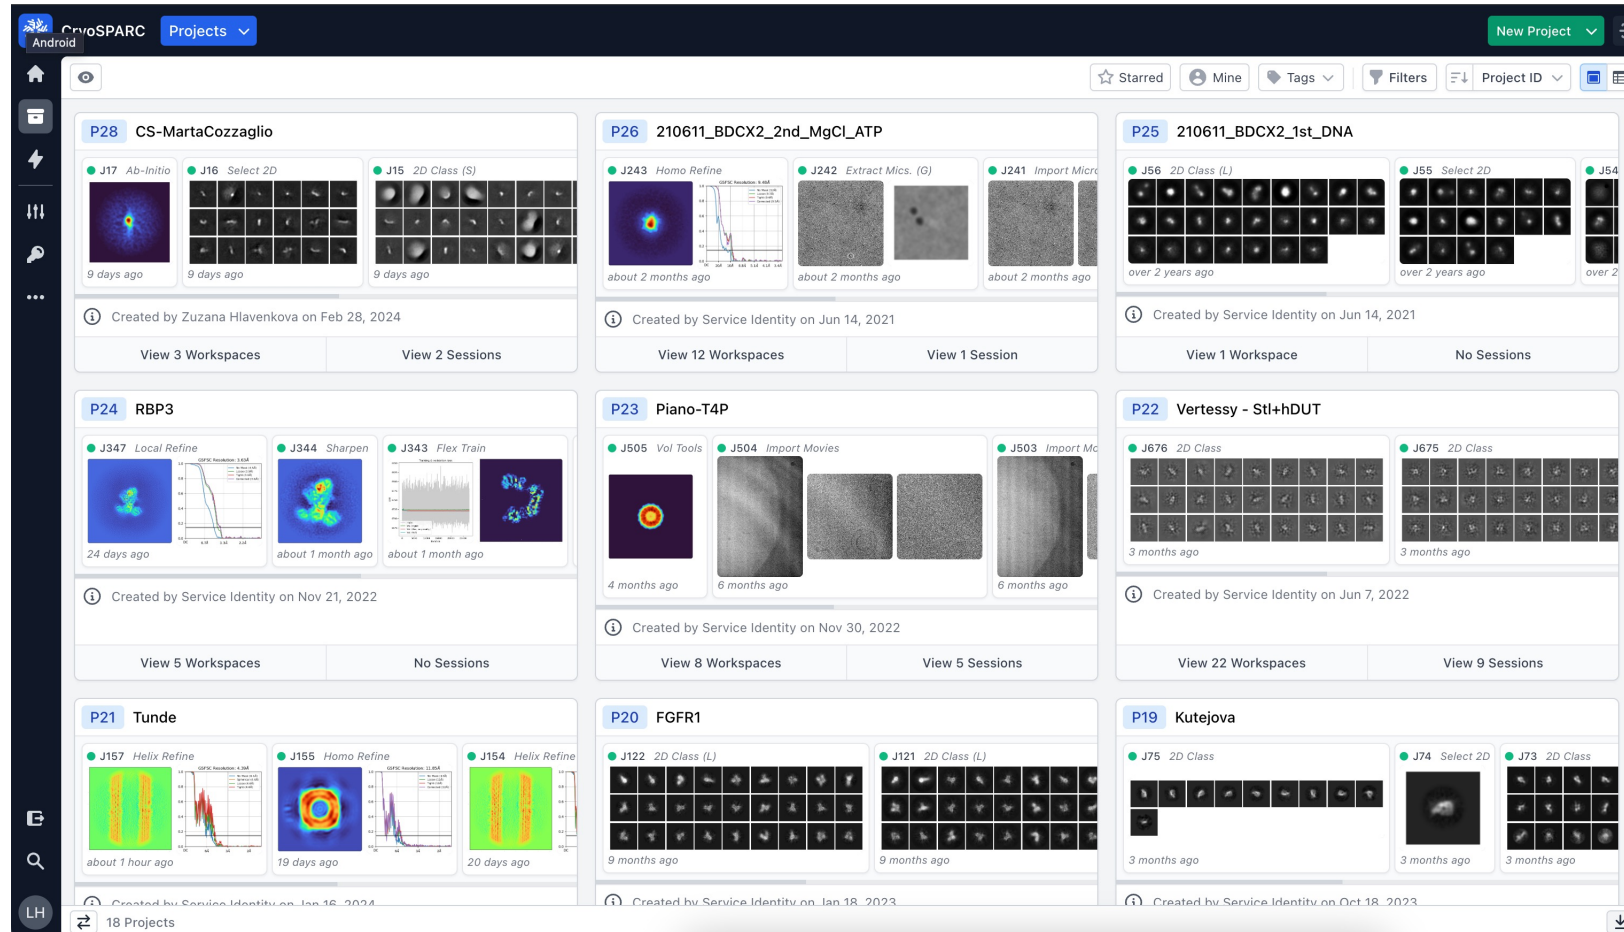


- Very low (pLDDT < 50)
- Low (70 > pLDDT > 50)
- Confident (90 > pLDDT > 70)
- Very high (pLDDT > 90)

Cryosparc

- Web UI running in K8s
- Jobs can run in both K8s and PBS
- Easy deployment of new instances
 - Own licence required

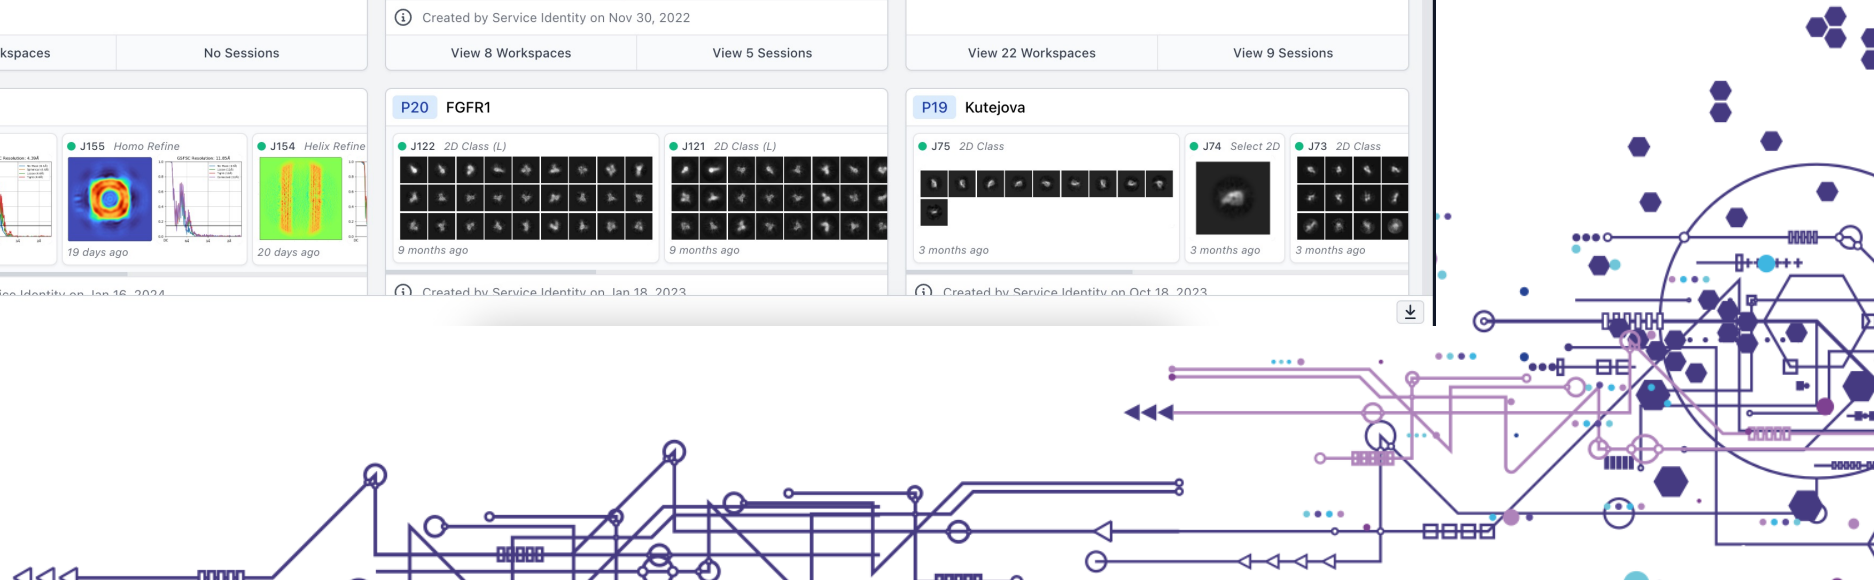




The screenshot displays the CryoSPARC web interface, showing a grid of project cards. Each card represents a project and contains a grid of thumbnails for workspaces and sessions. The projects shown are:

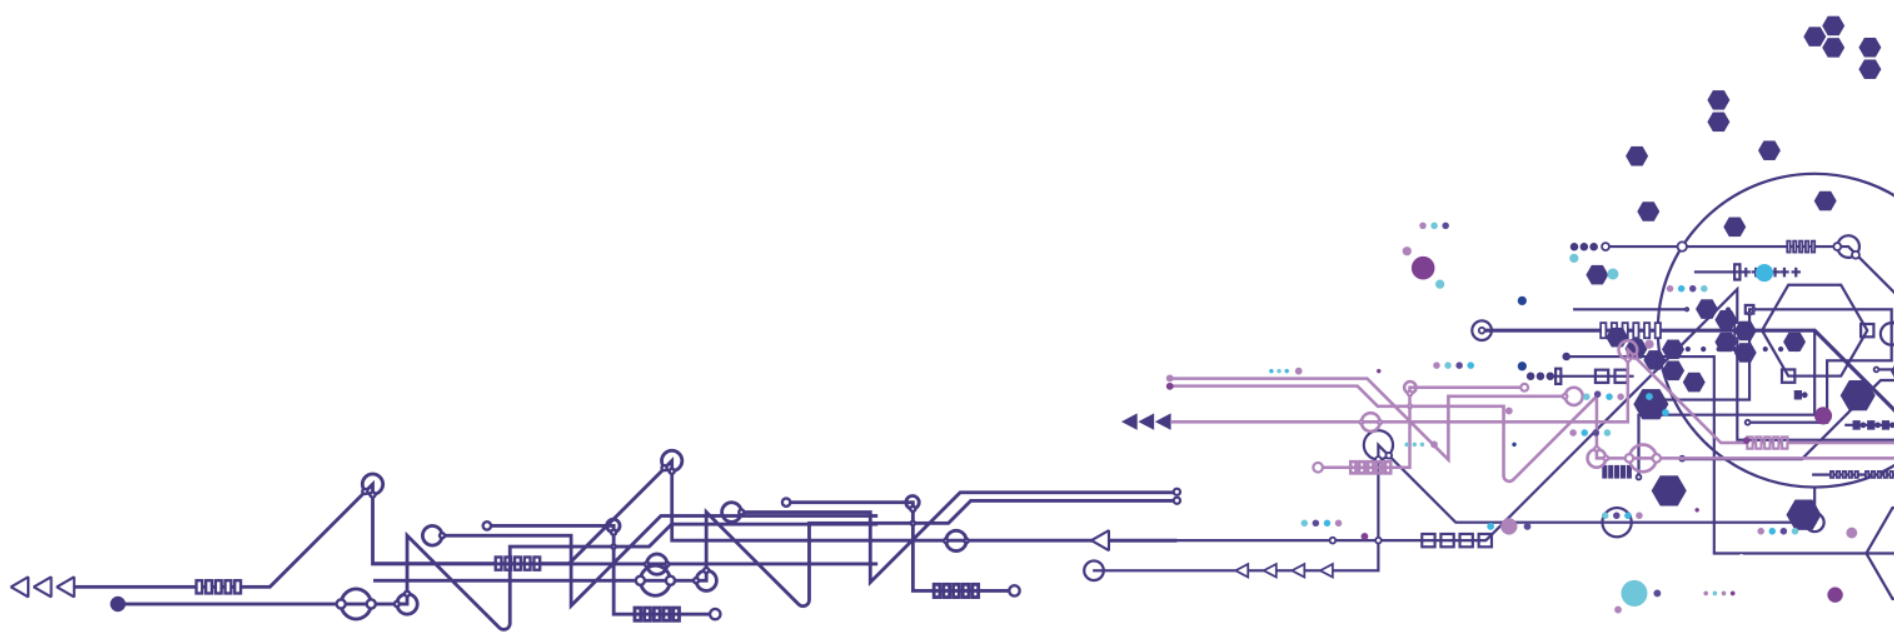
- P28 CS-MartaCozzaglio**: Created by Zuzana Hlavenkova on Feb 28, 2024. View 3 Workspaces, View 2 Sessions.
- P26 210611_BDCX2_2nd_MgCl_ATP**: Created by Service Identity on Jun 14, 2021. View 12 Workspaces, View 1 Session.
- P25 210611_BDCX2_1st_DNA**: Created by Service Identity on Jun 14, 2021. View 1 Workspace, No Sessions.
- P24 RBP3**: Created by Service Identity on Nov 21, 2022. View 5 Workspaces, No Sessions.
- P23 Piano-T4P**: Created by Service Identity on Nov 30, 2022. View 8 Workspaces, View 5 Sessions.
- P22 Vertessy - Stl+hDUT**: Created by Service Identity on Jun 7, 2022. View 22 Workspaces, View 9 Sessions.
- P21 Tunde**: Created by Service Identity on Jan 16, 2024. View 3 Workspaces, View 3 Sessions.
- P20 FGFR1**: Created by Service Identity on Jan 18, 2023. View 2 Workspaces, View 2 Sessions.
- P19 Kutejova**: Created by Service Identity on Oct 18, 2023. View 1 Workspace, View 3 Sessions.

The interface includes a top navigation bar with 'CryoSPARC Android', 'Projects', and 'New Project'. A sidebar on the left contains navigation icons. The top right of the grid has filters for 'Starred', 'Mine', 'Tags', 'Filters', and 'Project ID'.



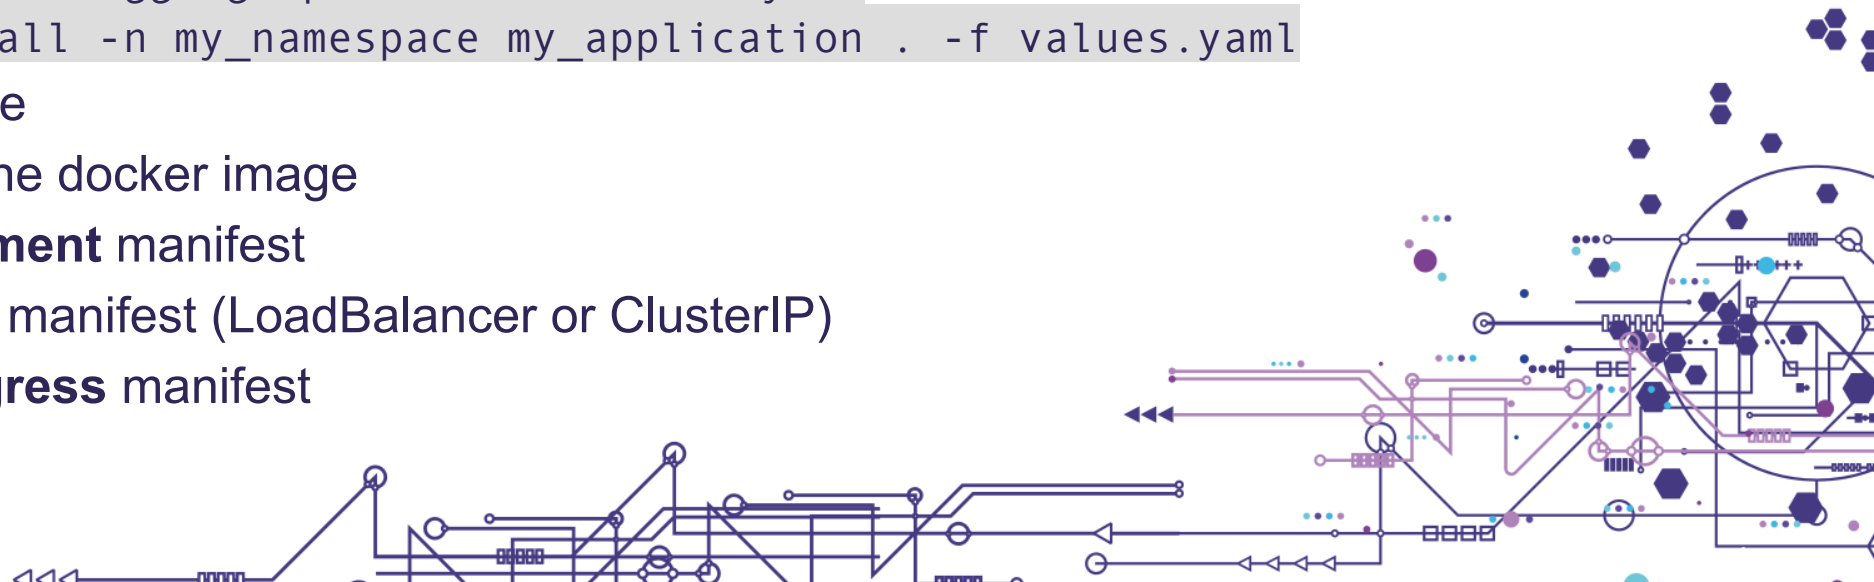


Native Kubernetes API



Native Kubernetes API Use Cases

- Get kubeconfig from Rancher
- Deploy **helm** package
 - Pre-packaged application to run in Kubernetes
 - Form of a template
 - Need to fill custom values
 - Examples:
 - ```
helm install --namespace logging logging-operator oci://ghcr.io/kube-logging/helm-charts/logging-operator -f values.yaml
```
    - ```
helm install -n my_namespace my_application . -f values.yaml
```
- Deploy **Docker** image
 - Find or create the docker image
 - Create **Deployment** manifest
 - Create **Service** manifest (LoadBalancer or ClusterIP)
 - Create **Ingress** manifest



Native Kubernetes API — Tips

- Exposing Web application to the Internet
 - Deployment + Service + Ingress
 - Let's encrypt certificate can be requested
- Exposing raw application to the Internet
 - Deployment + Service (LoadBalancer type)
 - Exposed at particular IP, can get ad-hoc name [.dyn.cloud.e-infra.cz](https://dyn.cloud.e-infra.cz)
 - Let's encrypt certificate can be requested
 - Can be used for, e.g., opening ssh connection to the container



Native Kubernetes API — Tips

- How to run Docker image that expects to be run as root
 - Mostly root privileges are not necessary
 - Mount `emptyDir` volume where write access is needed (e.g. logs)
 - Rebuild image
 - ```
FROM original_image AS src
```

```
FROM scratch
```


```
COPY --chown=1000:1000 --from=src / /
```

```
USER 1000
```
    - Does not work with Kaniko builder
    - Does not inherit `ENV`, `WORKDIR`, `CMD` settings
- How to allow additional packages to be installed?
  - Rebuild image as above
  - Install `fakeroot` (for deb based distros)
  - ```
ln -s /usr/bin/fakeroot /usr/bin/sudo
```



Native Kubernetes API — Common Problems

- would violate PodSecurity "restricted:latest": allowPrivilegeEscalation != false (container "demo" must set securityContext.allowPrivilegeEscalation=false), runAsNonRoot != true (pod or container "demo" must set securityContext.runAsNonRoot=true)
 - Deployment is missing proper securityContext
 - <https://docs.cerit.io/docs/securitycontext.html>
- Error: container has runAsNonRoot and image will run as root
 - Deployment or image is missing run user definition
 - Deployment: runAsUser: 1000
 - Image: USER 1000





Questions?
k8s@ics.muni.cz

A logo consisting of a large circle with a thick black outline. The text 'e-infra.cz' is centered within the circle. Two curved lines, one above and one below the circle, suggest motion or a partial ring.

e-infra.cz