e-INFRA CZ compute services

Adrian Rosinec, CESNET

IBERGRID 2022, Faro

cesnet
MUNI CERIT-SC
VŠB TECHNICKÁ UNIVERZITA OSTRAVA | IT4INNOVATIONS NÁRODNÍ SUPERPOČÍTAČOVÉ CENTRUM

Compute services overview

- Grid based on PBS
- OpenStack IaaS cloud
- SensitiveCloud - PaaS
- Managed Kubernetes – PaaS
- Supercomputer

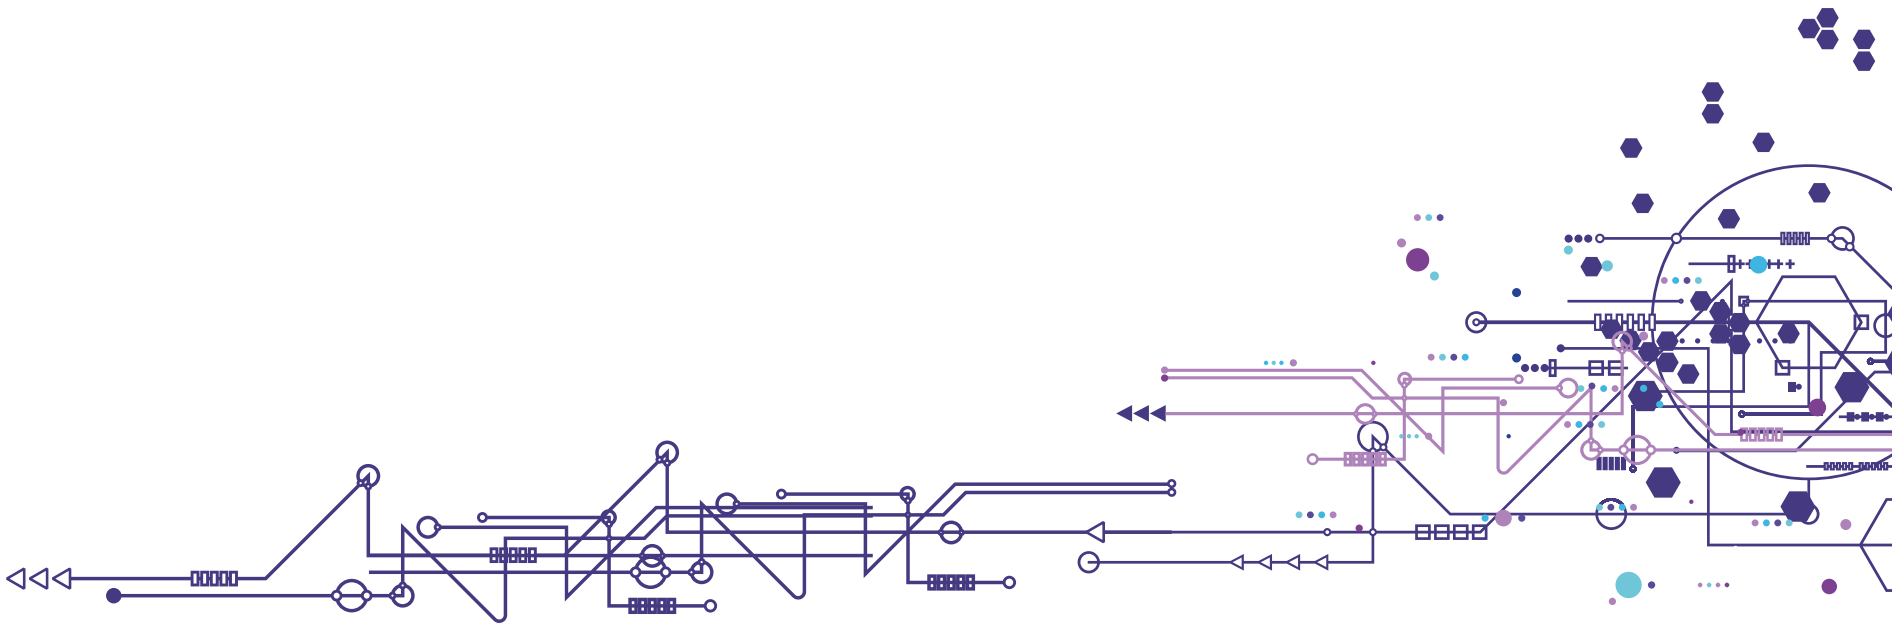

Cloud Services

Current topics in cloud services

Cloud services as tool to support research

e-INFRA CZ

Advanced Computing for Research

elixir CZECH REPUBLIC

RECETOX

cesnet metacentrum

eosc

MUNI ICS

C4e

BBMRI-ERIC®

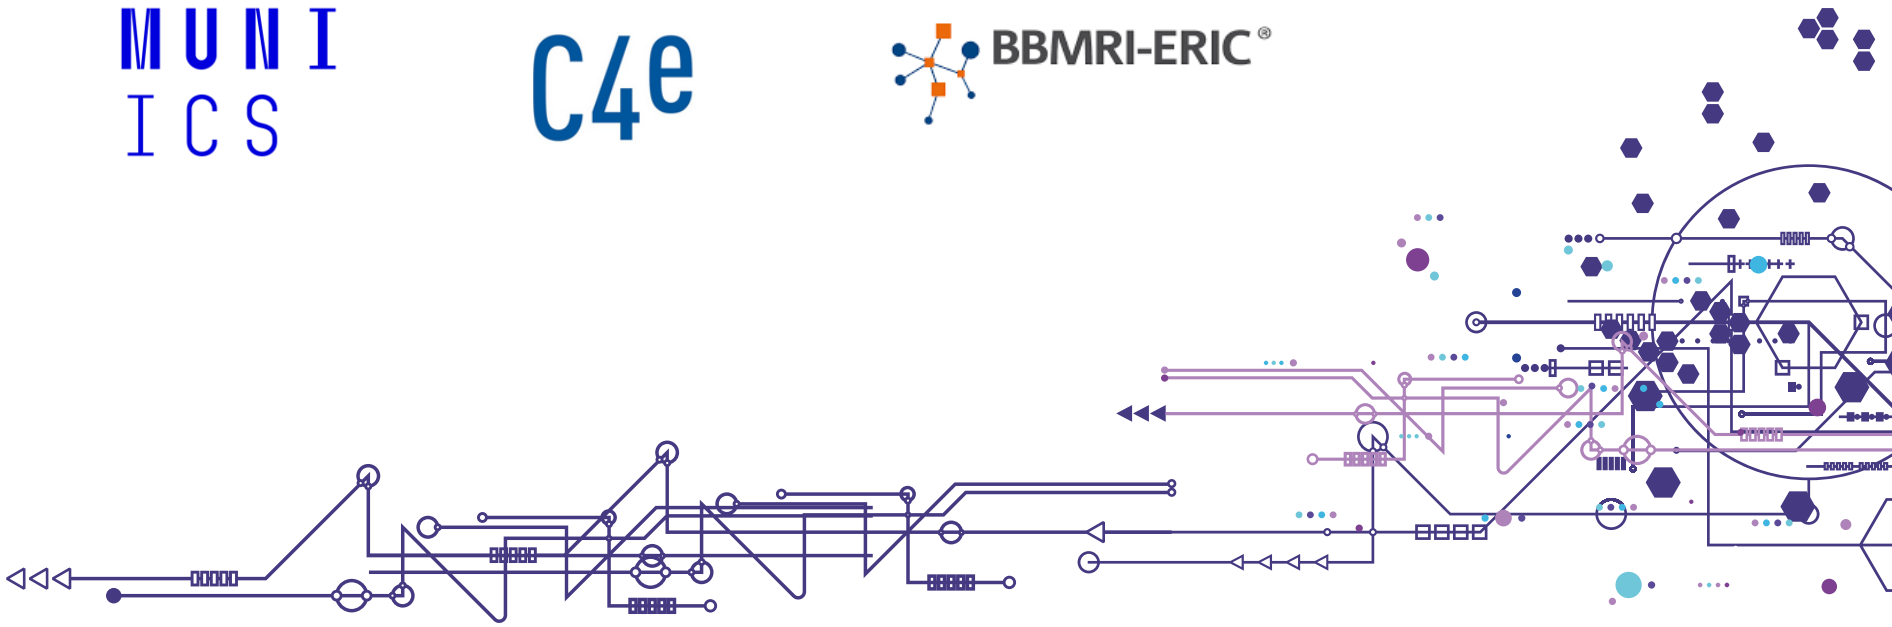

Sensitive Data Processing

Services built to process sensitive data

- Introduction of SensitiveCloud BETA
  - ISO 27k certified compute and storage solution
  - Secured private network perimeter – access through VPN (pki-based Wireguard)
  - 2FA to management
- Sensitive Compute built on top of container platform
  - “focus only on app maintenance = better overall security”
- Sensitive Storage - block storage
  - Easy integration to Kubernetes via NFS-CSI
- Strict network policies, network multi-tenancy
- Driven by Bio/Health use-cases and stakeholders
- Need for IaaS

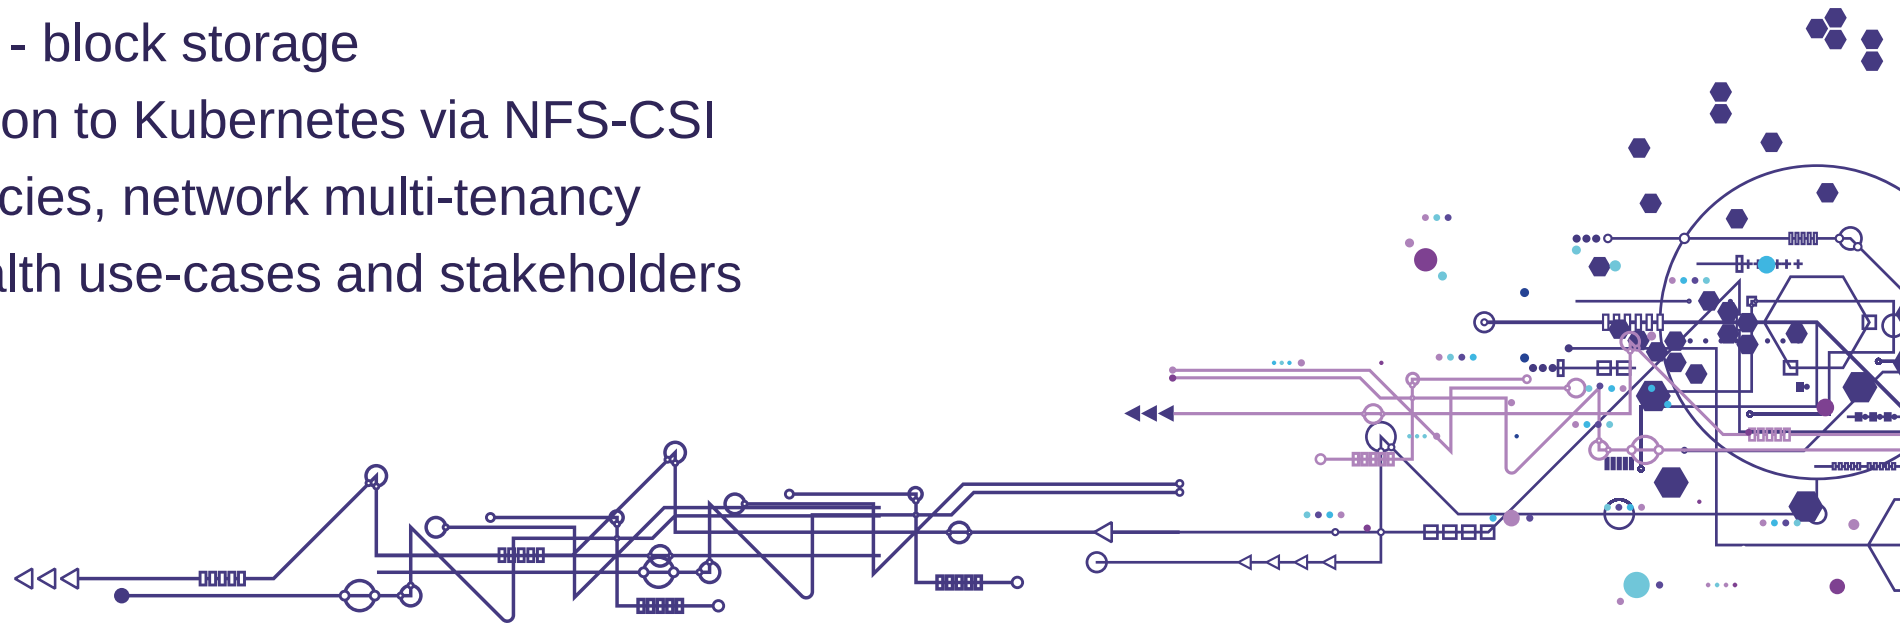

OpenStack IaaS cloud

- 9000 vCPU, 200TB of RAM, 10PB storage
- Optimized for HPC flavors
  - 100 GPUs – A40, T4, RTX also available to EGI communities
  - Fast ephemeral NVME storage
  - InfiniBand
- Virtual networking, CEPH object/block, LBaaS, Auto-Image rotation
- 200 national research & education activities using IaaS
- 32 EGI use-cases - *biomed, pangeo, panosc, c-scale, icecube, carousel,* …
- Variety of users
  - Automatization maturity – terraform/ansible vs. dashboard
  - HPC/HTC jobs and semi-production services (web, …)

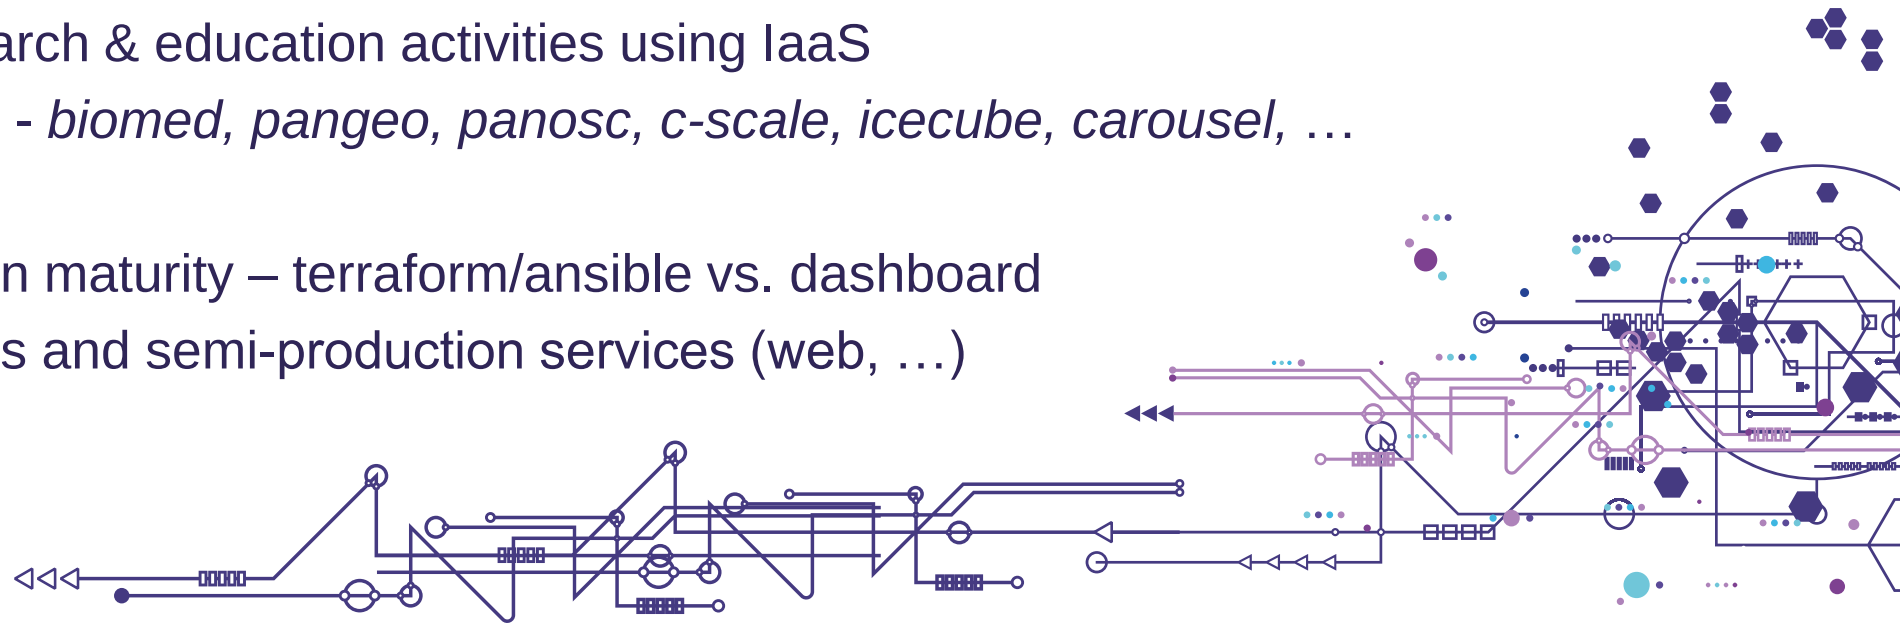

EOL OpenStack gen I.

Brace for OpenStack gen II., built on the new platform

- Since 2017 based on OpenStack-Kolla
- Non sustainable own orchestrating management Puppet-Kolla
- Ended up in the mix of Ansible+Puppet, Kolla containers on control plane
- Hard to maintain and upgrade OpenStack to newer versions
  - Tracking and resolving incidents
  - Applying changes
  - Expensive further development

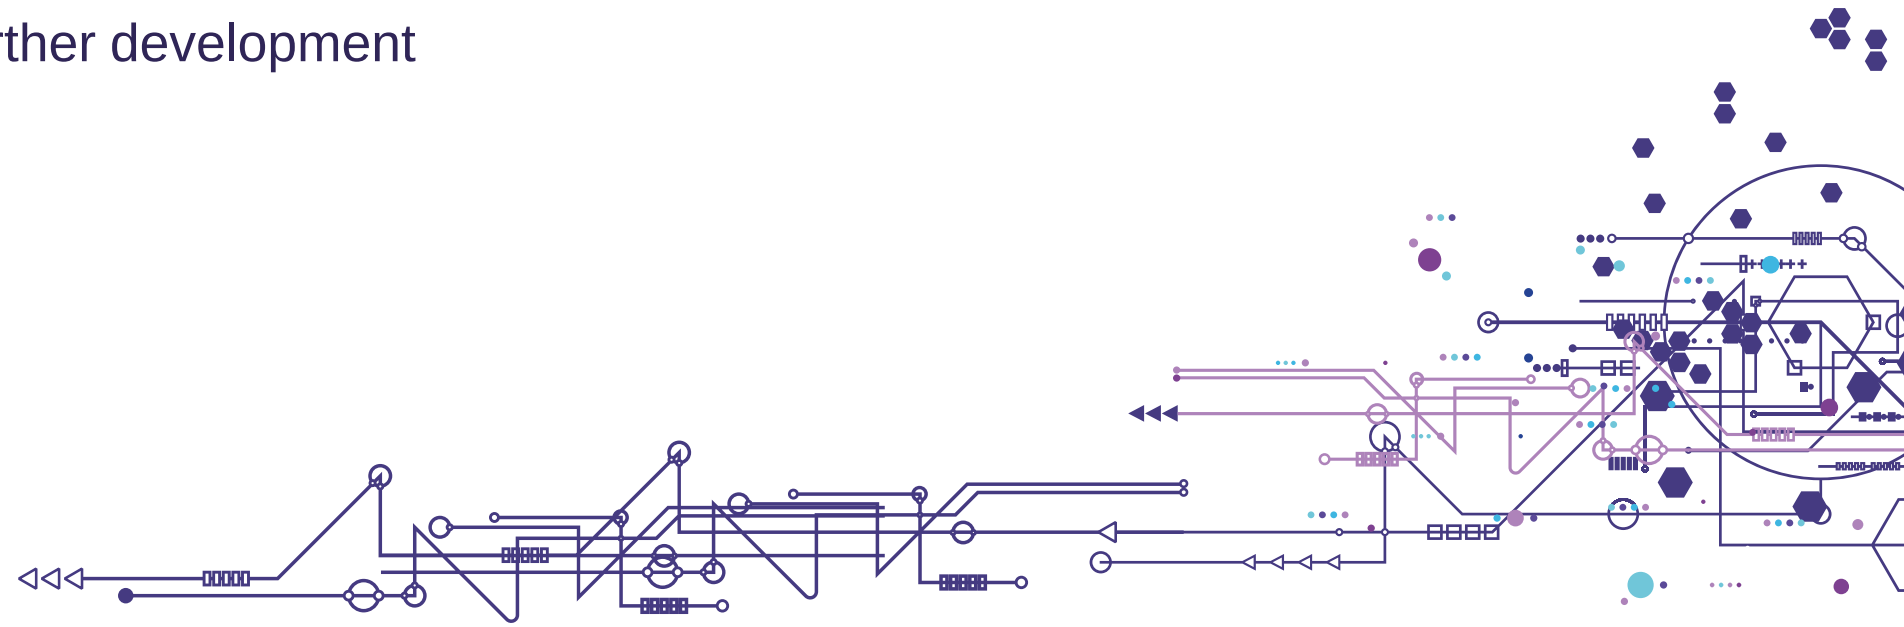

EOL OpenStack gen I.

New architecture

- Hardware installation - MaaS
- Kubernetes as orchestrator of control plane
- OpenStack components as HELM charts
- Utilization of *GitOps* and *Infrastructure as Code*
  - administration not done on “admin” machines with their local scripts
  - whole intelligence in pipelines
  - also great SPoF
- Requirements (metrics) of new architecture
  - ops:dev and use-case engineering – 20:80
- Resulting OpenStack distribution will be open-source project
- New hardware – all flash CEPH, more GPUs, (virtual GPUs)

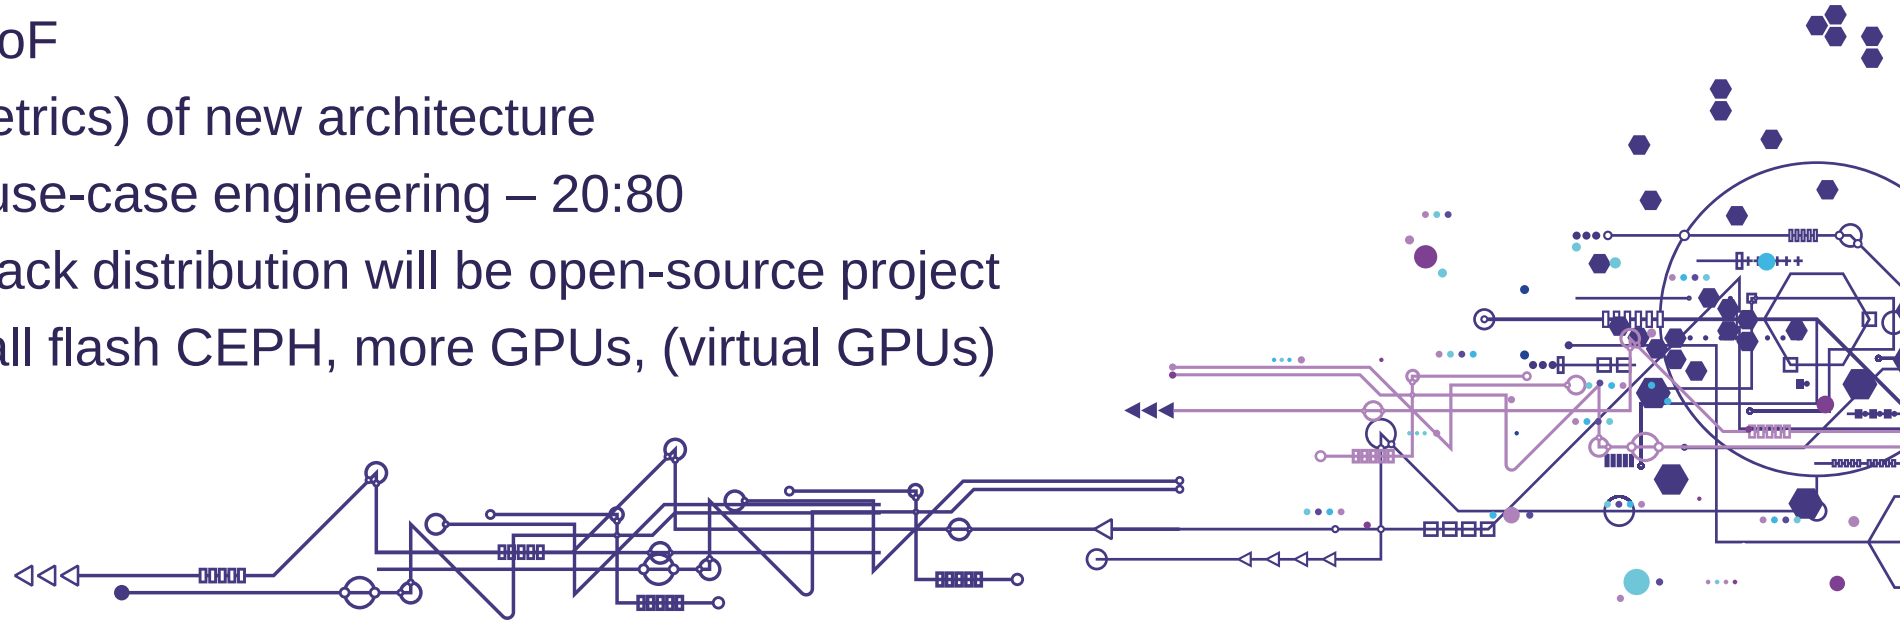

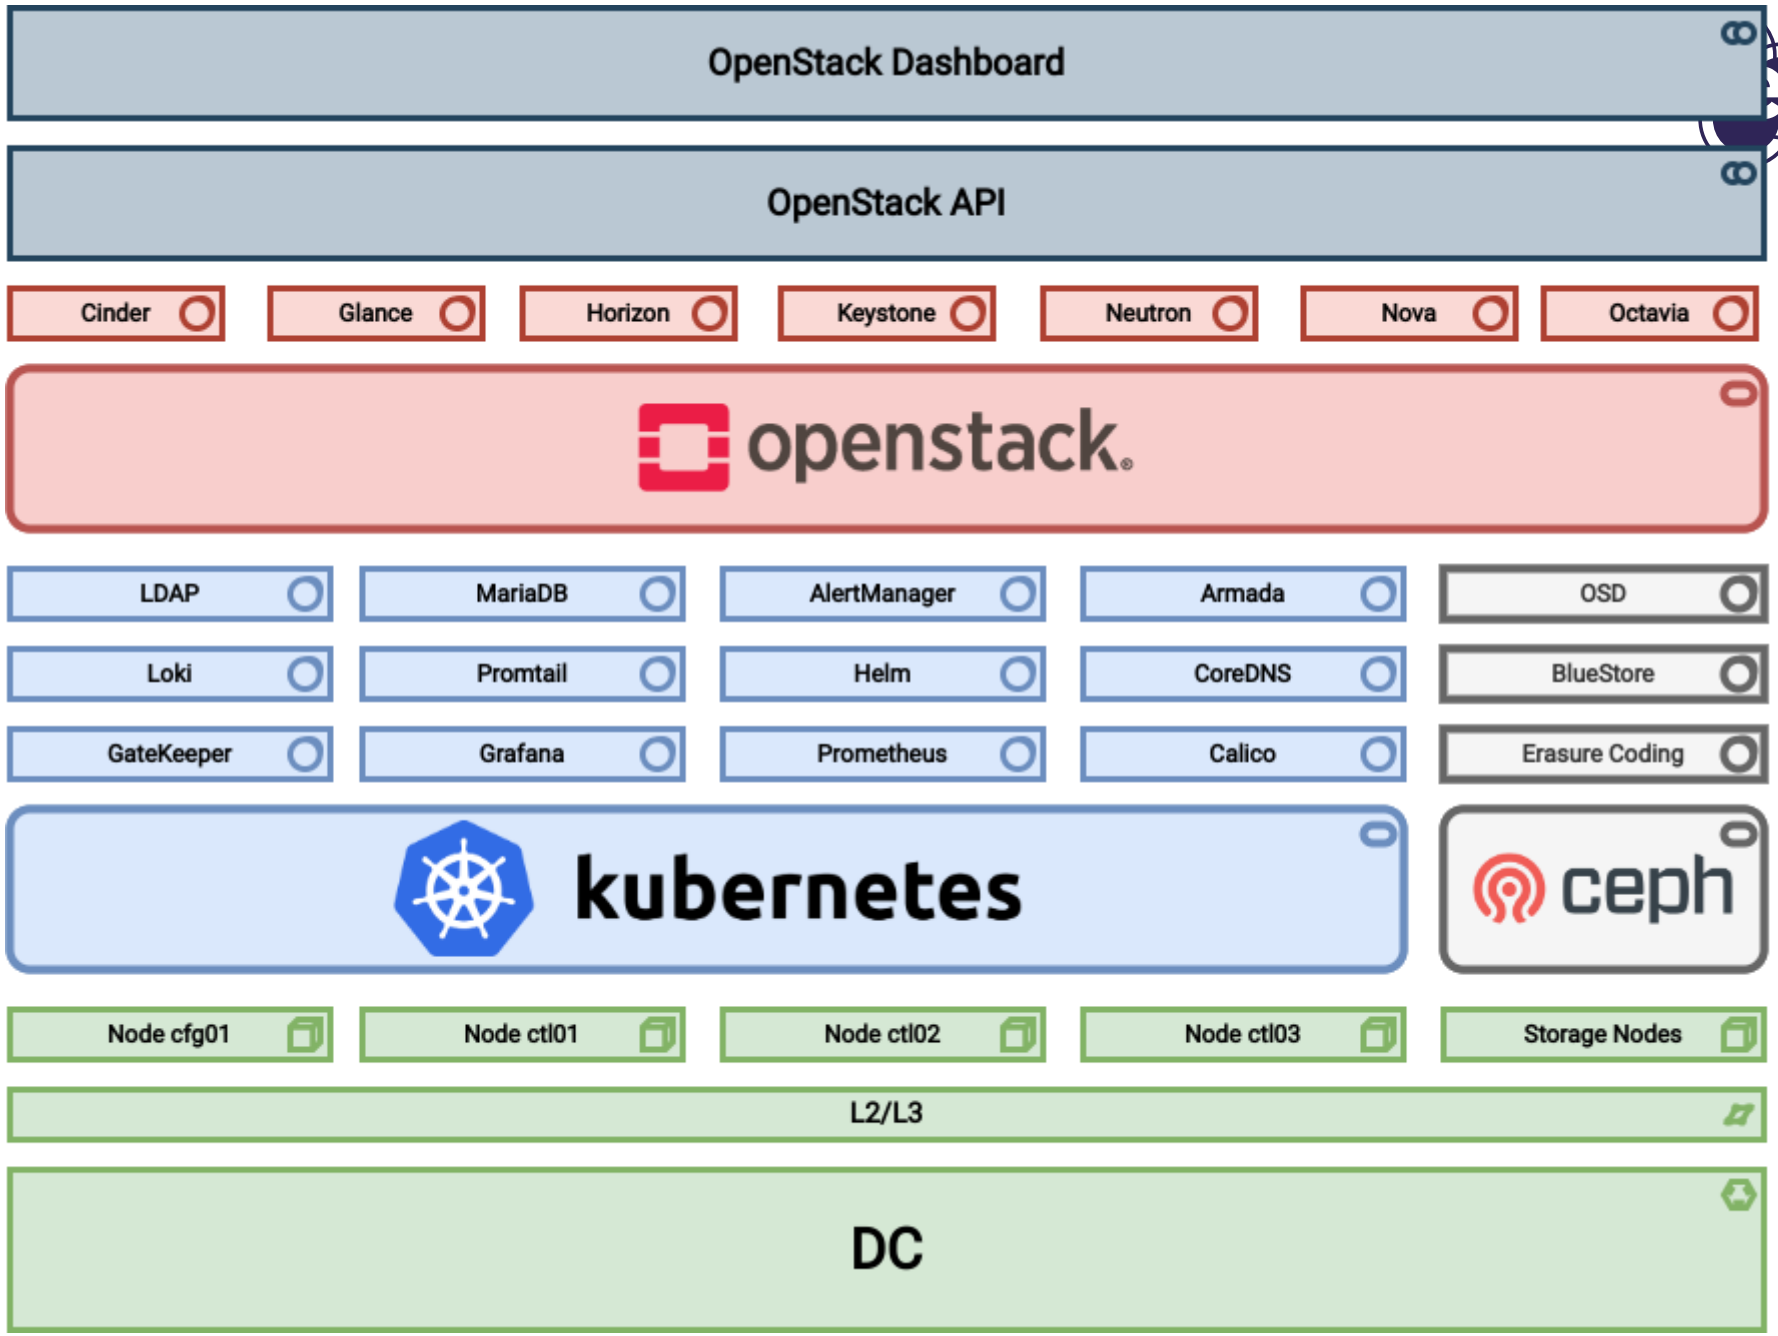
OpenStack Dashboard
OpenStack API
Cinder
Glance
Horizon
Keystone
Neutron
Nova
Octavia
openstack.
LDAP
MariaDB
AlertManager
Armada
OSD
Loki
Promtail
Helm
CoreDNS
BlueStore
GateKeeper
Grafana
Prometheus
Calico
Erasure Coding
kubernetes
ceph
Node cfg01
Node ctl01
Node ctl02
Node ctl03
Storage Nodes
L2/L3
DC

Container platform

- Containers are getting popular in scientific community
- We want our users to focus on research activities rather than maintenance/ops (dev vs. ops)
- Managed shared pool of resources, Kubernetes orchestrator = managed container platform
- Managed DNS/CertMan, Ingress, PVs, GPUs
- Supporting svcs – Container registry, CI/CD, HELM repo, … (working on with EGI)
- Use-cases
  - Nextflow pipelines
  - Microservice apps
  - SaaS apps (interactive Matlab + license, Jupyter Notebooks, Desktops, …)
- Research in domain of k8s workload scheduling – resource fairness