Converging HPC, Big Data and Cloud technologies for precision agriculture data analytics on supercomputers

Yiannis Georgiou

ryax technologies

Naweiluo Zhou, Li Zhong, Dennis Hoppe - **HLRS**
Marcin Pospieszny - **PSNC**
Nikela Papadopoulou, Kostis Nikas, Orestis Lagkas Nikolos - **ICCS**
Pavlos Kranas - **LEANXSCALE**
Sophia Karagiorgou - **UBITECH**
Eric Pascolo - **CINECA**
Michael Mercier, Pedro Velho - **RYAX**

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

CYBELE

VHPC'20

V-camp

HPC, Big Data and Cloud technologies

- **HPC** traditionally used for scientific computing leveraging supercomputers parallel and distribute computing techniques
- **Big Data** used for data analytics to extract valuable insights utilizing **Cloud** data centers with elastic environments

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

V-camp

Digital Transformation in Agritech

- **Big Data analytics** can be applied in various industries including **agriculture** and **farming**
- **HPC** can benefit Big Data applications since large datasets can be processed in timely manner and with improved accuracy
- Combining Big Data and HPC can enable **precision agriculture** and **precision livestock farming**

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

V-camp

HPC complexity–Big Data/HPC Convergence

- **HPC** steep learning curve
- Traditional Big Data/Cloud **containerization** and **orchestration** cannot be used out-of-the box in HPC due to security and performance issues.
- Workflows Combining Big Data and HPC executions still a research topic

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

Outline

- Introduction – CYBELE project context
- Architecture for HPC abstraction
- Big Data - HPC meta-scheduling with Kubernetes
- Environment deployment on HPC through Singularity
- Multi-GPU scaling Precision Agriculture Experiments
- Conclusions – Perspectives

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

V-camp

Introduction – CYBELE project

- **CYBELE** is an **EU-funded** project aiming to revolutionize agriculture, aquaculture and livestock farming by combining:
    - **High Performance Computing**
    - **Big Data / AI Analytics**
    - **Cloud Computing**

https://www.cybele-project.eu/

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

V-camp

Introduction – CYBELE project

- **4 HPC Testbeds**

PSNC

HLRS
High-Performance Computing Center | Stuttgart

Bull
atos technologies

CINECA

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

V-camp

Introduction – CYBELE project

- **9 Agriculture/Farming Pilots**

1)Organic Soya Yield & Protein Prediction
2)Grape Production & Harvesting Models
3)Organic Fruit Production
4)Robotic systems for Arable Frameworks
5)Crop Yield Forecasting
6)Pig Weighing Optimisation
7)Sustainable Pig Production
8)Open Sea Fishing
9)Aquaculture monitoring and feeding

Goal – CYBELE project

- **4 HPC Testbeds**

- **9 Agriculture/Farming Pilots**

1)Organic Soya Yield & Protein Prediction
2)Grape Production & Harvesting Models
3)Organic Fruit Production
4)Robotic systems for Arable Frameworks
5)Crop Yield Forecasting
6)Pig Weighing Optimisation
7)Sustainable Pig Production
8)Open Sea Fishing
9)Aquaculture monitoring and feeding

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

Goal – CYBELE project

- **4 HPC Testbeds**

PSNC

H L R

High-Performance Computing

Bull

atos technologies

CINECA

- **9 Agriculture/Farming Pilots**

1)Organic Soya Yield & Protein Prediction
2)Grape Production & Harvesting Models
3)Organic Fruit Production

Make HPC more accessible and improve AgriTech data analytics speed and accuracy

7)Sustainable Pig Production

8)Open Sea Fishing

9)Aquaculture monitoring and feeding

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

V-camp

Paper Context and Contributions

- **Typical** production HPC systems provide **static combination** of Big Data and HPC partitions
- We propose a prototype architecture with support of **hybrid BD-HPC executions** where Kubernetes manages a VM-based Big Data partition while Slurm/Torque manage baremetal HPC partitions
- Containerization is possible through **Docker or Singularity** on **Big Data partition** and only through **Singularity** on **HPC partitions**
- **Data Analytics** applications can be executed **on HPC partitions** using **Kubernetes API**, through a meta-scheduling technique **based on wlm-operator, singularity-CRI, multi-user support and Kubernetes communication with Slurm/Torque**
- **Environment Deployment techniques** based on **Singularity** and repository with **pre-built images** for Big Data and AI frameworks

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

Architecture for HPC Abstraction

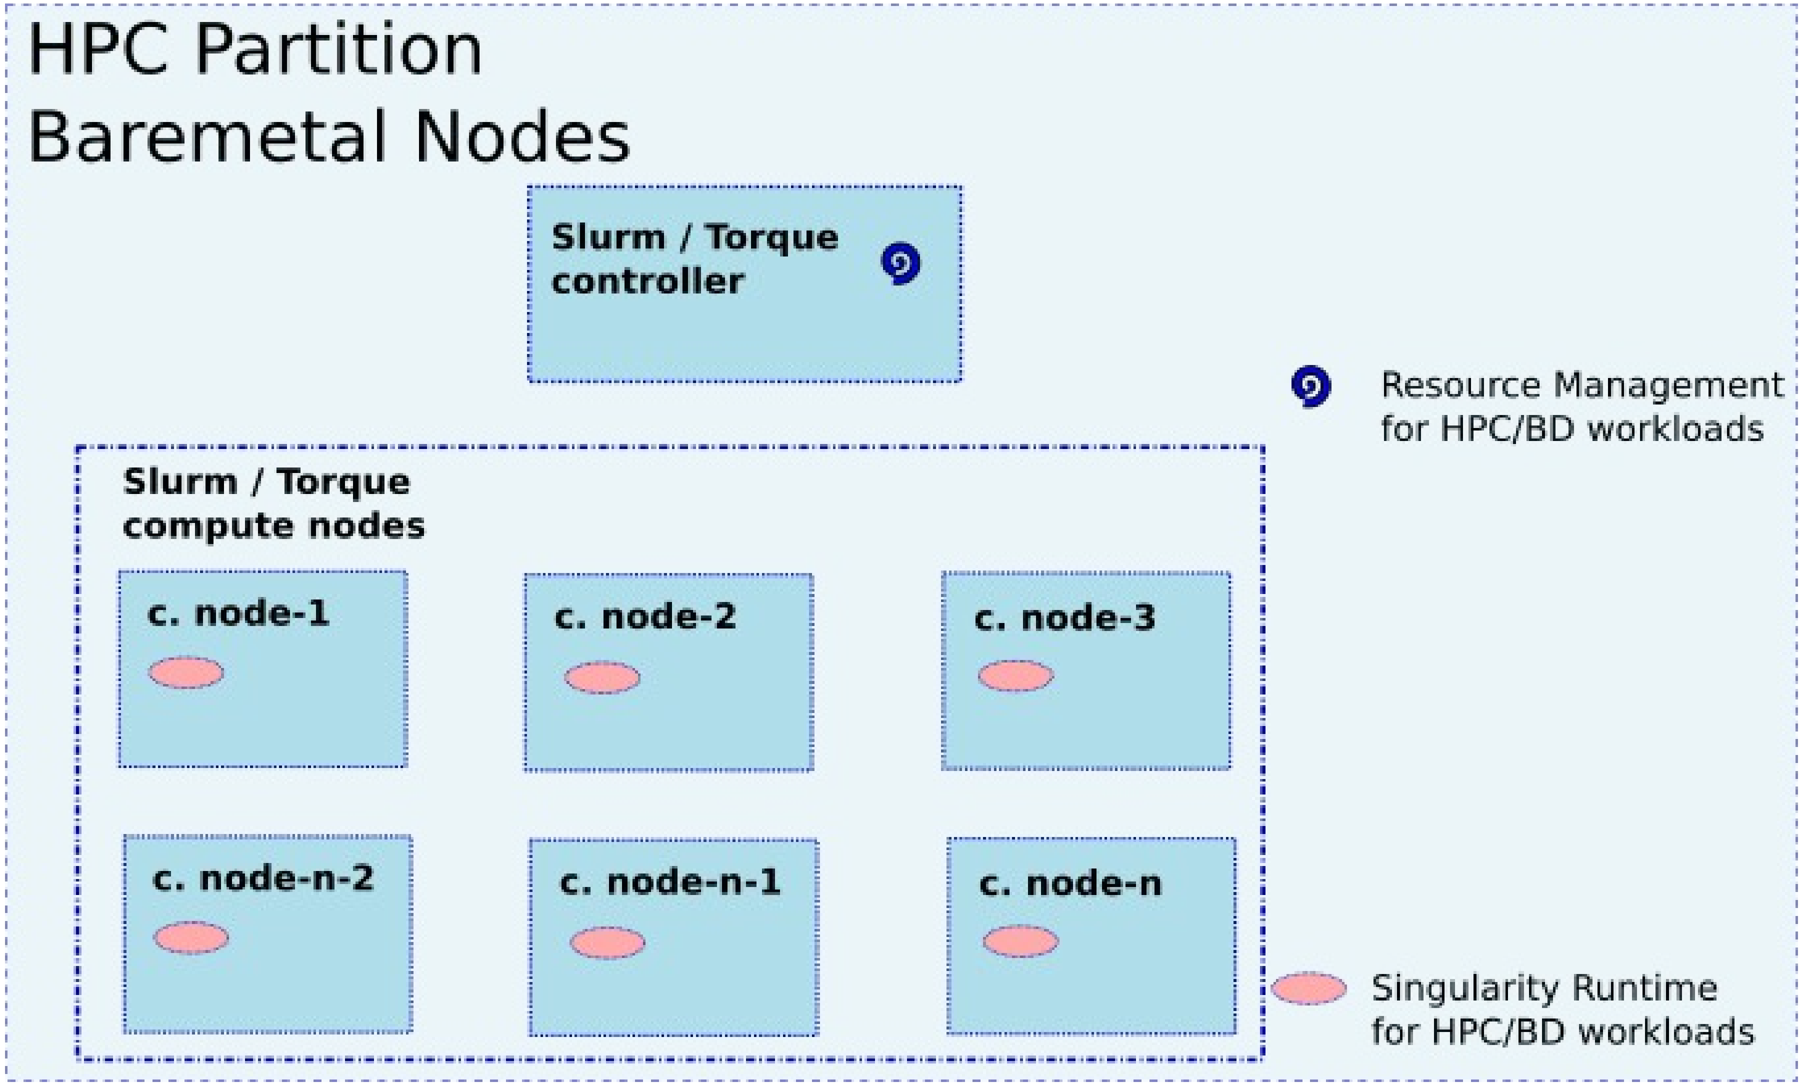

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

V-camp

Architecture for HPC Abstraction

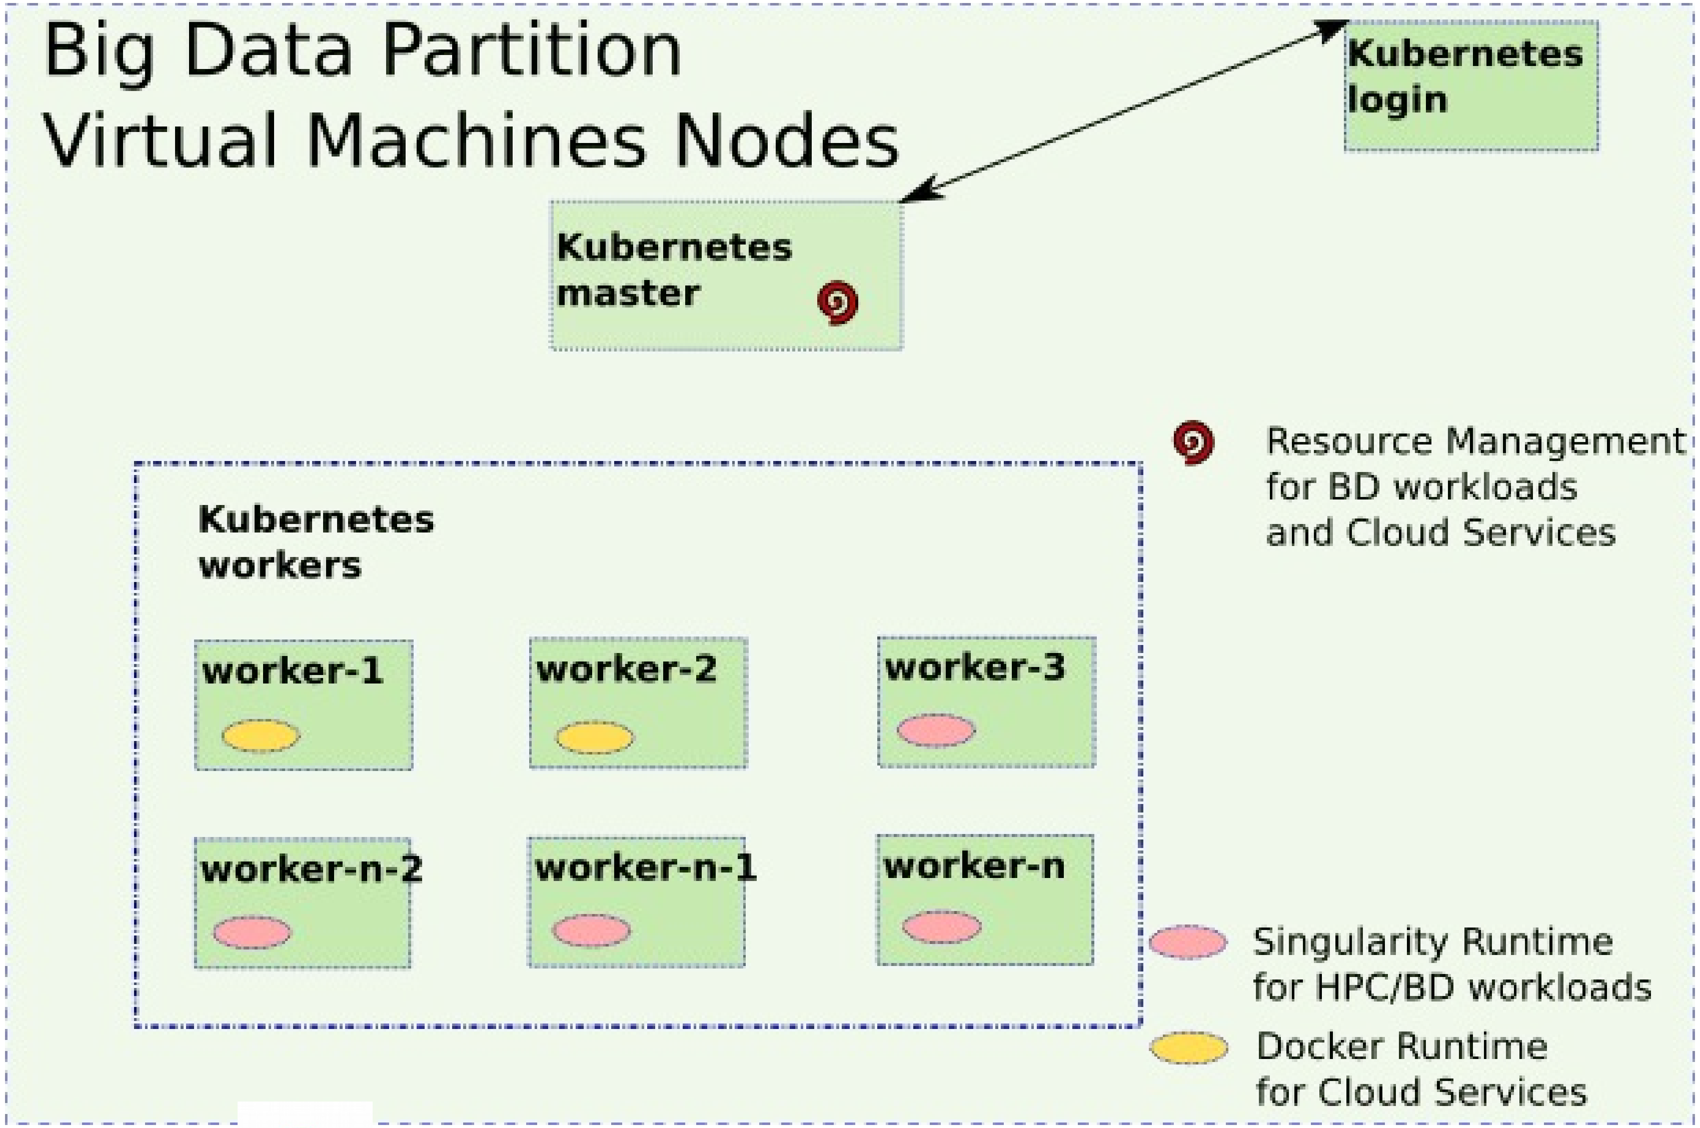

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

Architecture for HPC Abstraction

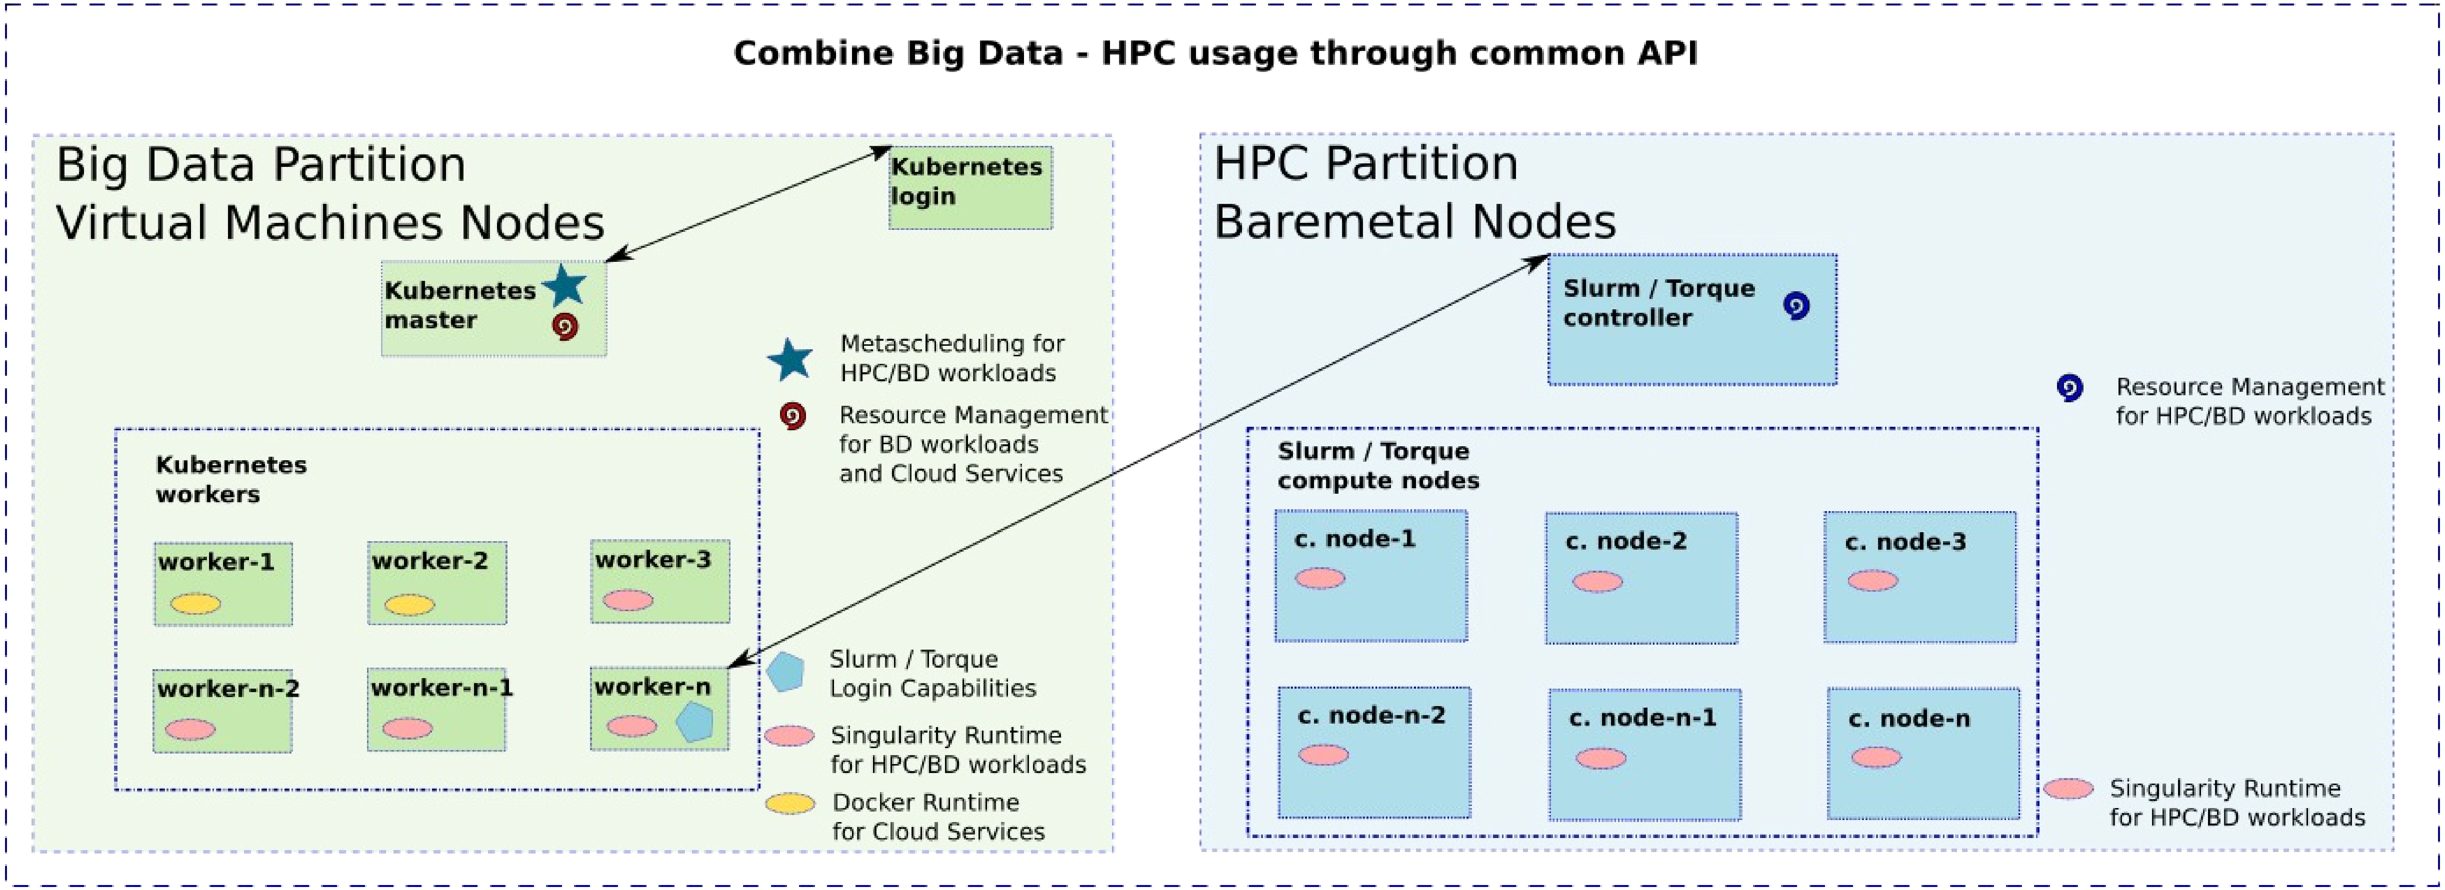

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

V-camp

Architecture for HPC Abstraction

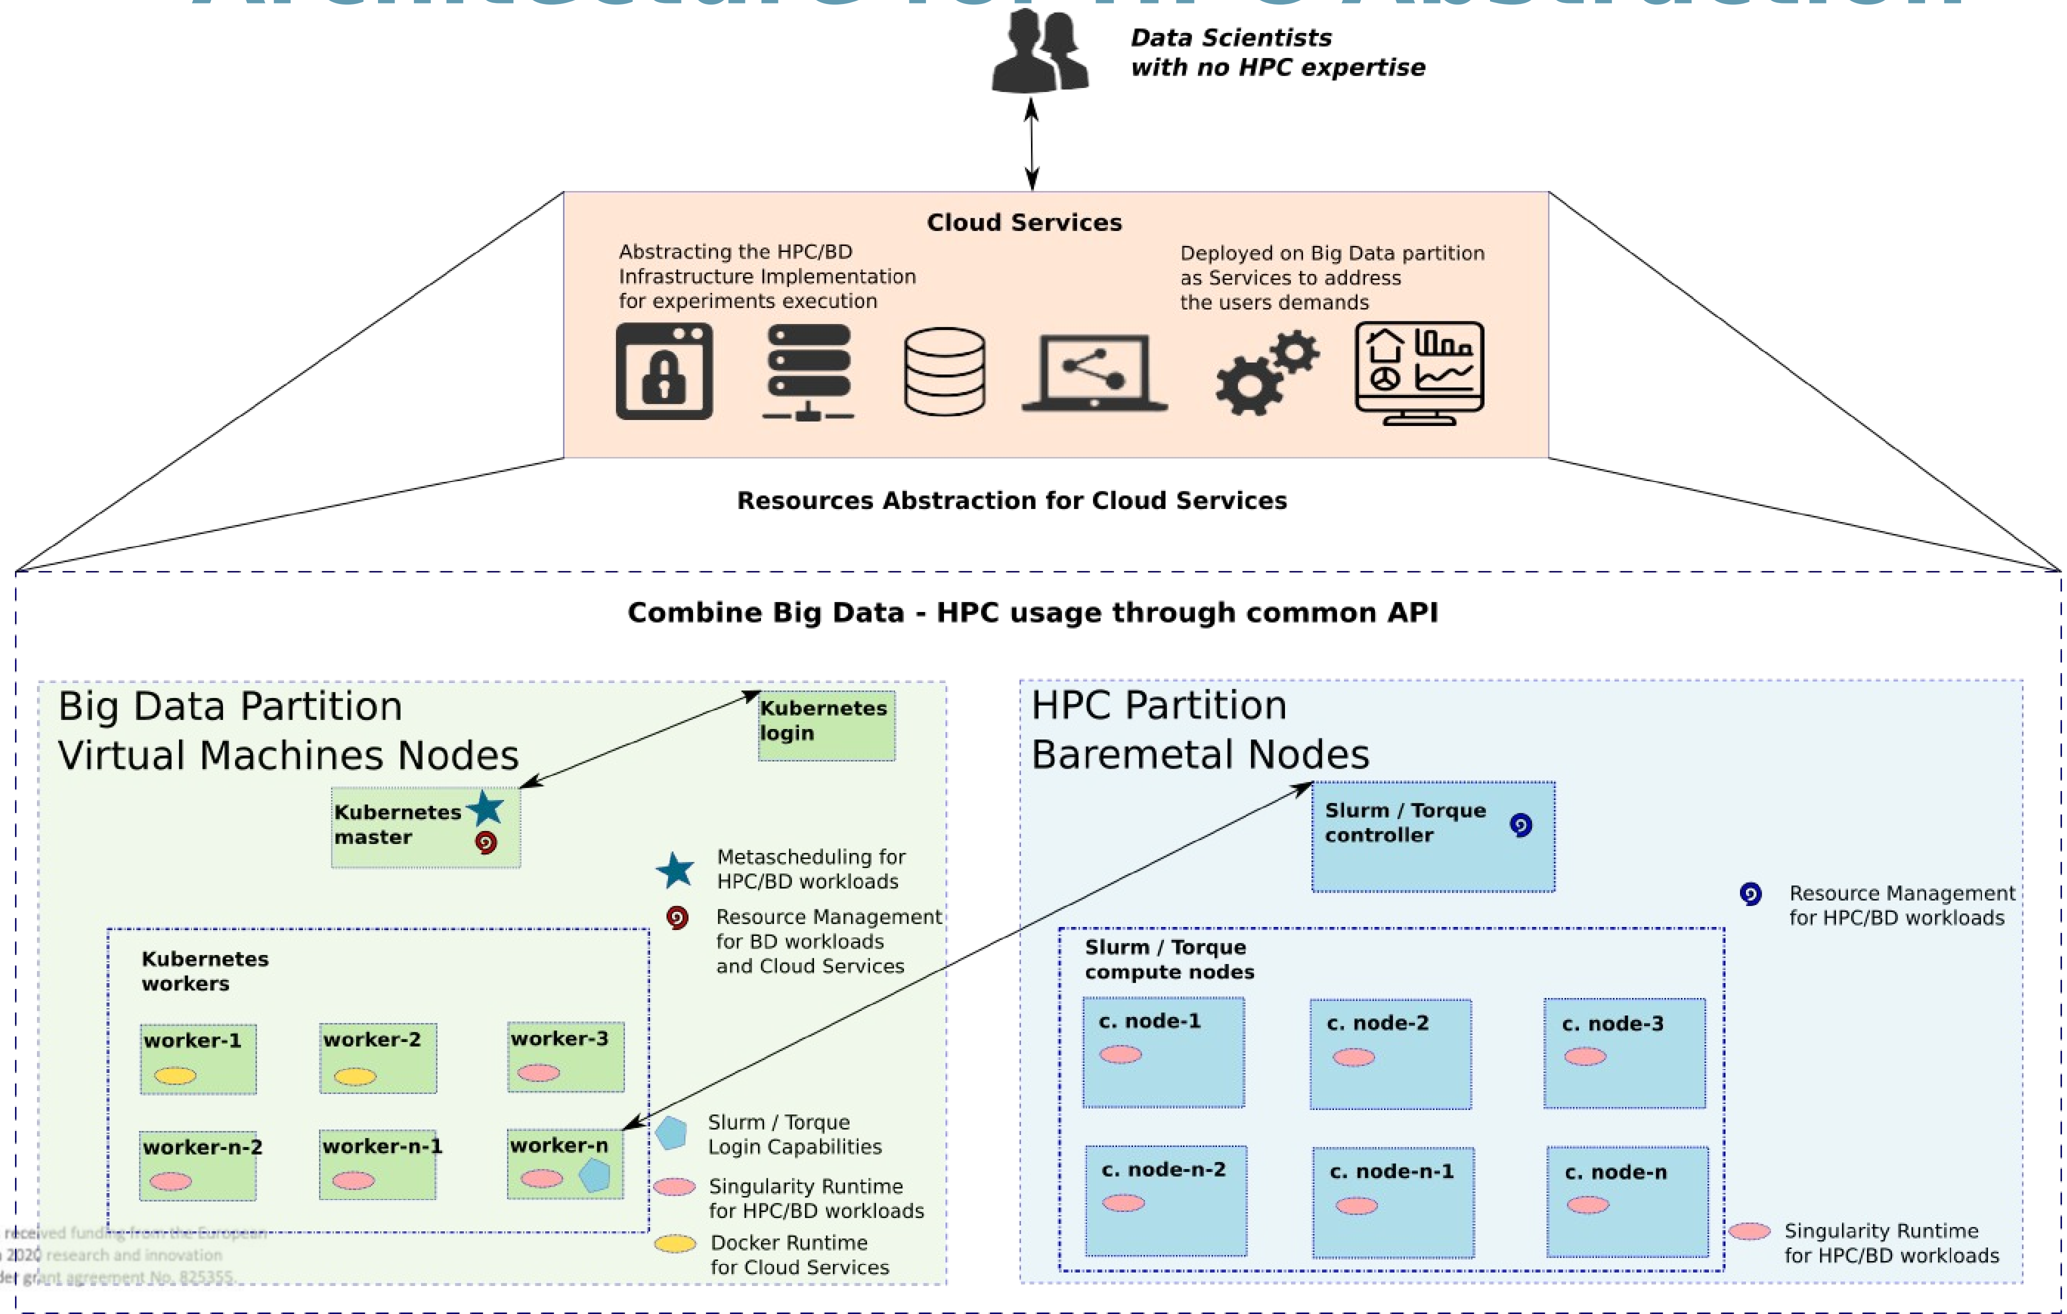

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

Cloud-BD–HPC meta-scheduling with K8S

- Kubernetes orchestration for **adapted runtime usage** for Cloud services and Data Analytics on the Big Data partition :
  - **Docker CRI** for Cloud Services (Workflow Management, Databases, Visualization, etc)
  - **Singularity CRI**[1] for Data Analytics applications
  - Scheduling based on **labels** (for workers) and **node-selectors** (for executions/pods)
- **Common Kubernetes API** for Data Analytics execution on both **Big Data and HPC partitions**:
  - Based on **wlm-operator software**[2] for communication between K8S and Slurm/Torque
  - K8S worker with Singularity CRI and Slurm/Torque login capabilities
  - Automated creation of virtual K8S workers (**virtual-kubelet**) per partition/queue
  - Communication protocol based on a **gRPC proxy per user** from the specific K8S worker node
  - **Multi-user support** through dynamic adaptation of the user context by automatically reconfiguring the virtual-kubelet and the necessary gRPC proxy to enable execution isolation monitoring, accounting, etc.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

[1]: https://github.com/sylabs/singularity-cri
[2]: https://github.com/sylabs/wlm-operator

Cloud-BD–HPC meta-scheduling with K8S

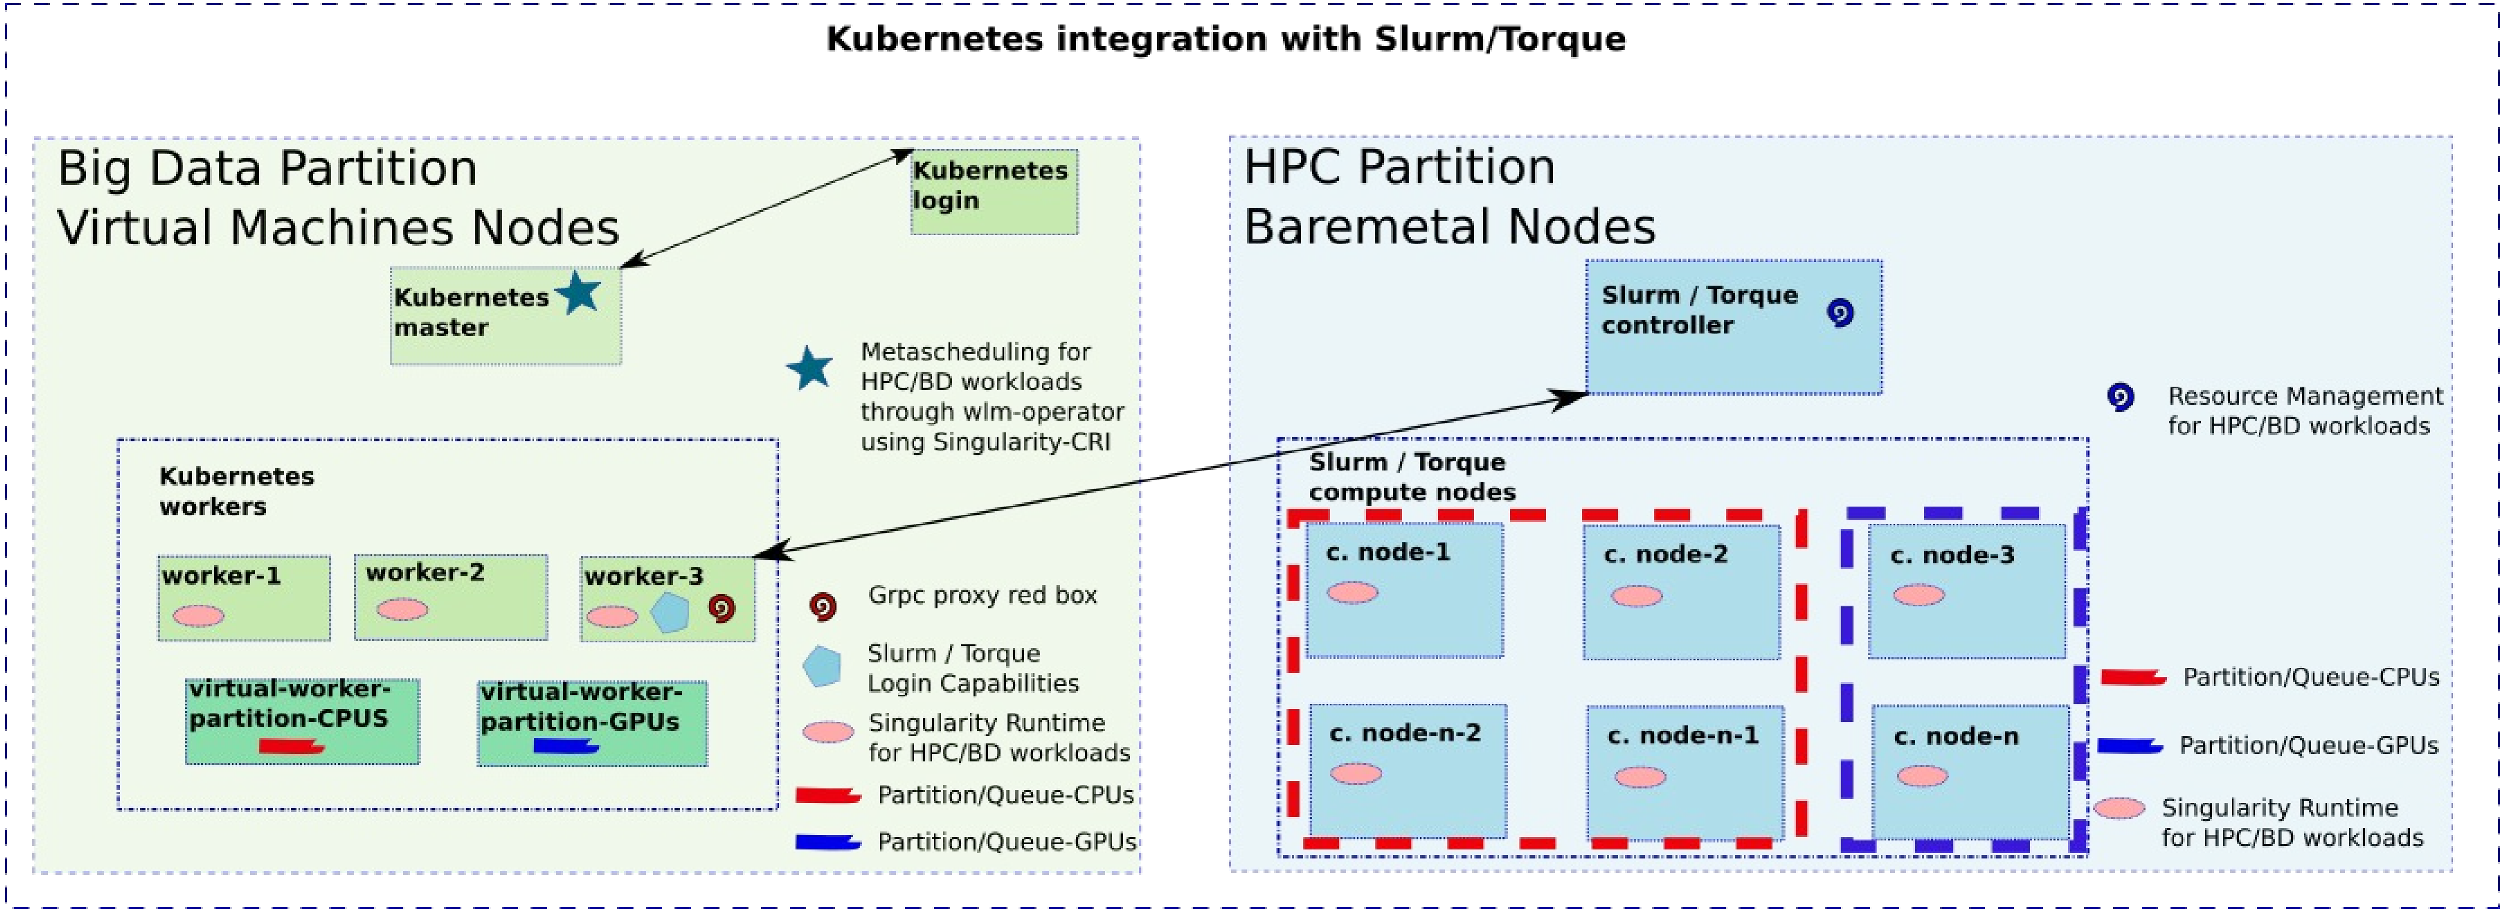

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

V-camp

Cloud-BD–HPC meta-scheduling with K8S

```
$cat slurm-batch-example.yaml
apiVersion: wlm.sylabs.io/v1alpha1
kind: SlurmJob
metadata:
  name: Slurm-batch-example
spec:
  batch: |
    #!/bin/sh
    #SBATCH --nodes=3
    #SBATCH --partition=CYBELE
    srun sleep 5
    srun hostname
...
$ kubectl apply -f batch_example.yaml
$ kubectl get pods
NAME                      READY    STATUS    RESTARTS  AGE
slurm-batch-example-job  0/1      Running       0       4s
…
$ squeue
JOBID    PARTITION   NAME     USER   ST TIME NODES
```

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

Multi-GPU Scaling for Precision Agriculture

- Performed some experiments using one real-life **precision agriculture application**.
- The aim of the application is to develop a framework for **automatic identification and counting of wheat ears** in fields by getting data from sensors on ground that will enable crop yield prediction at early stages and provide more informed decisions for sales planning.
- The application consists in **training a deep learning algorithm** written in **Python** and using **Fastai/Pytorch framework** based on a group of **RGB images** (initially 138 images).
- In particular we deployed the wheat ears counting application upon one single BD or HPC node testing the scaling and parallelization of the code by increasing the number of GPUs.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

Multi-GPU Scaling for Precision Agriculture

- The experiments have been performed on a **dedicated testbed** where the previously described architecture of Kubernetes orchestration on both Big Data and HPC partitions has been deployed, along with the integration to Slurm and Singularity for the execution on the HPC partition.
- The HPC testbed is part of **BULL NOVA cluster** and we made use of the following hardware:
  - **one HPC BareMetal node**, featuring a **Bull Sequana S800 machine**, equipped with 4X2 Intel Xeon Platinum 8253 (256CPUs), 4 TB RAM and 4 GPUs NVIDIA GV100GL Tesla V100 PCIe 16GB,
  - **one Big Data VirtualMachine node**, with 4 CPUs and 8GB RAM

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

Multi-GPU Scaling for Precision Agriculture

VirtualMachine(VM) or BareMetal(BM)	VM	BM	BM	BM	BM	BM
number of CPUs	4	256	256	256	256	256
number of GPUs	0	0	1	2	3	4
Execution Time (sec)	37008	7020	417	312	274	247

Table 1. Execution time of wheat-ears application on 1 VirtualMachine (4 CPUs) node or one BareMetal (256 CPUs) node scaling from 0 to 4 GPUs

- Both cases use **containerization with Singularity** for the executions with the difference that in the second case it is done on VM while in the first on BareMetal.
- The results in table 1 show the **performance improvement** of our application when using a powerful **HPC BareMetal** node with GPU instead of **small VM**, **100* orders of magnitude.**
- Scaling of GPUs impacts the application performance: **10* orders of magnitude** when using **GPUs rather than only CPU.**

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

V-camp

Conclusions

- Proposed a prototype architecture to enable the **execution of Big Data Analytics** upon **supercomputers** using different Big Data and HPC hardware partitions and a **converged Big Data-HPC-Cloud software stack.**
- **Kubernetes** as high-level **orchestrator** and **common API** to allow the deployment of Data Analytics as HPC jobs.
- An **environment deployment tool** bringing pre-built Singularity images of **Big Data and AI frameworks** (Pytorch,Tensorflow, etc) specifically adapted to **targeted HPC resources** (GPUS, Infiniband, etc)
- Mechanisms to be used as **basic building blocks** to enable **HPC abstraction** and through the standard Kubernetes API to **deploy higher-level Cloud tools** to simplify the life of Data Scientists eliminating the need of learning new tools to make use of HPC platforms.
- Code to be made publicly available as open-source at the end of the project.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

Perspectives

- Propose different pre-built images for BD/AI frameworks to support various types of HPC hardware.
- Perform end-to-end experiments combining workflows with hybrid Big Data – HPC executions.
- **Further optimizations** will be done for the **multi-user support** of the Kubernetes integration with Slurm/Torque to be **more tightly integrated** with the communication mechanism.
  - **Adaptations** can be also done by exploring the new **Slurm REST-API** for more direct communications with Cloud services
- Work towards the support of **specialized Big Data frameworks** such as **Spark**
  - which can make use of **Kubernetes as resource manager** on Big Data partitions and
  - deploy **Spark applications** on **Slurm clusters** through a **non-interfering** method of **low-priority jobs** [3].

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825355.

[3] Michael Mercier, David Glesser, Yiannis Georgiou, Olivier Richard:
Big data and HPC collocation: Using HPC idle resources for Big Data analytics. BigData 2017: 347-352

Thanks!

yiannis.georgiou@ryax.tech

Yiannis Georgiou et al. Converging HPC, Big Data and Cloud Technologies for Precision Agriculture Data Analytics on Super computers. ISC Workshops 2020: 368-379

- https://www.cybele-project.eu/

https://www.ryax.tech