A deep dive on BigHPC platform delivery into HPC clusters

Speaker: Samuel Bernardo (LIP) on behalf of BigHPC consortium

BIG HPC HIGH PERFORMANCE COMPUTING

Cofinanciado por: COMPETE 2020 Lisb@20 20 PORTUGAL 2020 UNIÃO EUROPEIA Fundo Europeu de Desenvolvimento Regional FCT Fundação para a Ciência e a Tecnologia UTAustin Portugal

Outline

- The BigHPC consortium
- BigHPC platform overview
- Development challenges
- Integration Activities
- Takeaways

BigHPC consortium

Austin, USA

UT Austin

TACC

Portugal

Wavecom

INESC TEC

LIP

MACC

BigHPC platform overview

	Platform Global Services	HPC Local Applications
Storage Manager	Storage Global Controller	Storage Local Controller
Monitoring	MONICA	Hector
Virtualization Manager	VM Controller VM Repository	VM Executor

Monitoring backend - Overview

- Monitoring probe inside compute nodes responsible for collecting and pushing metrics
- Aggregator of metrics from several probes
- SQL based database with support for PromQL queries
- Visualisation dashboards

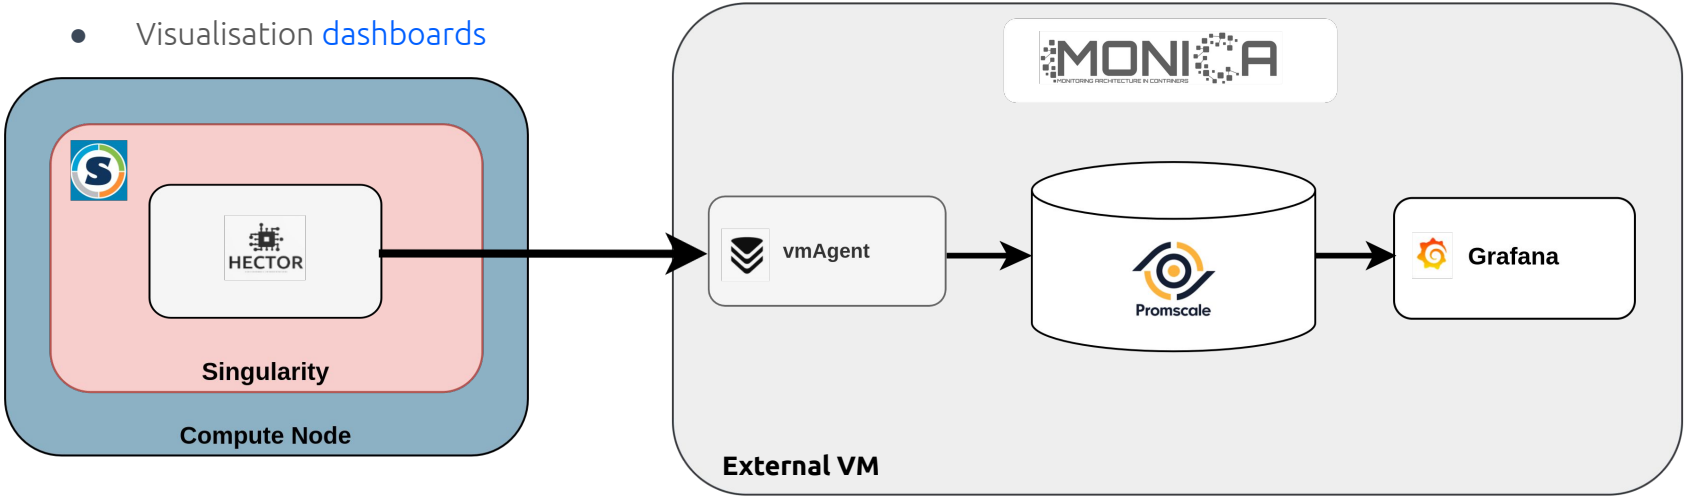

Virtualization Manager - Architecture

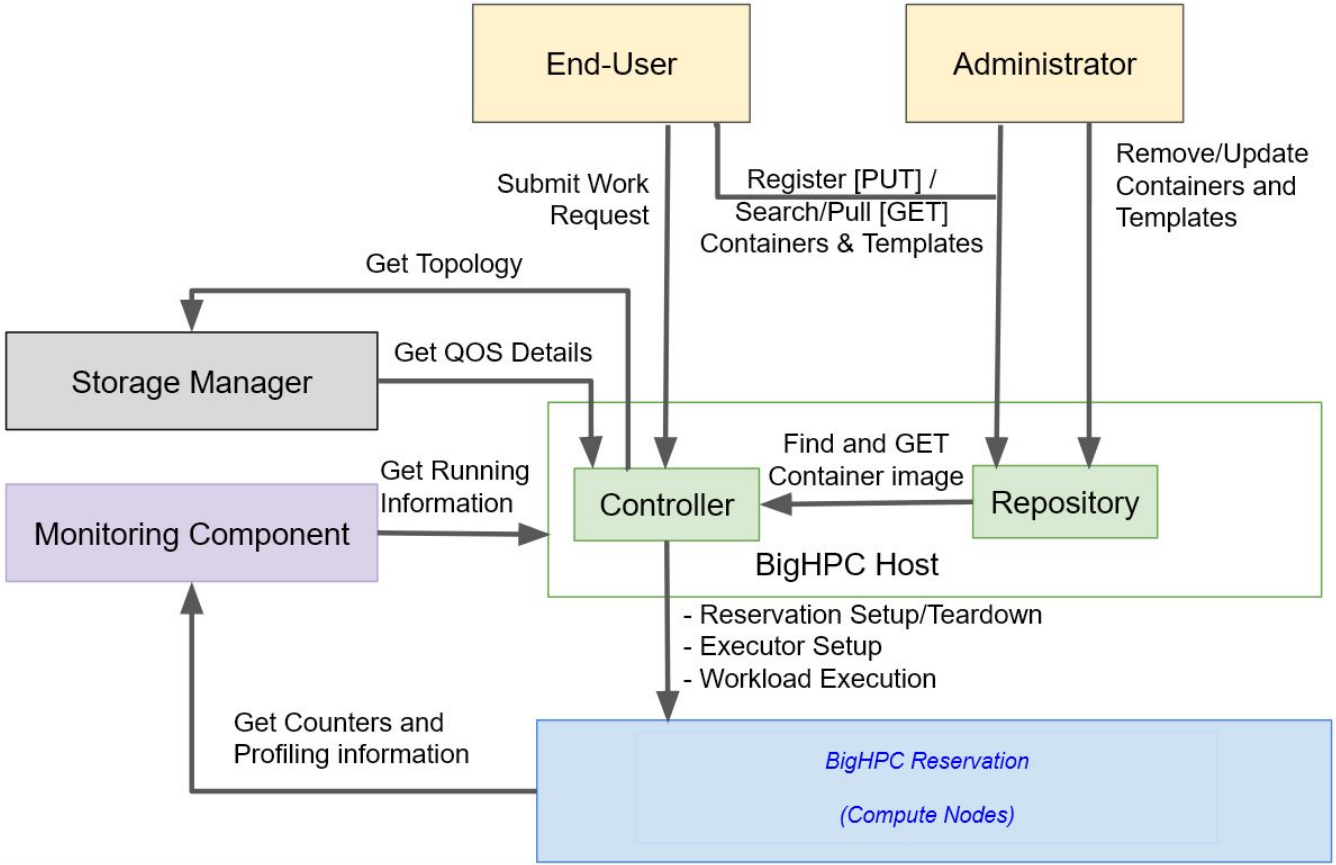

Software-Defined Storage - Overview

- Improve storage performance and management for HPC services
  - Ensure Quality of Service for I/O-intensive jobs
  - Alleviate I/O pressure at the shared parallel file system

- Key design features
  - Target both Big Data and Scientific workloads
  - Transparent to applications
  - Agnostic of the underlying storage resources (e.g., local and shared file systems)

- Based on the Software-Defined Storage paradigm
  - Data plane stages manage per application I/O requests
  - Control plane ensures holistic coordination across data plane stages

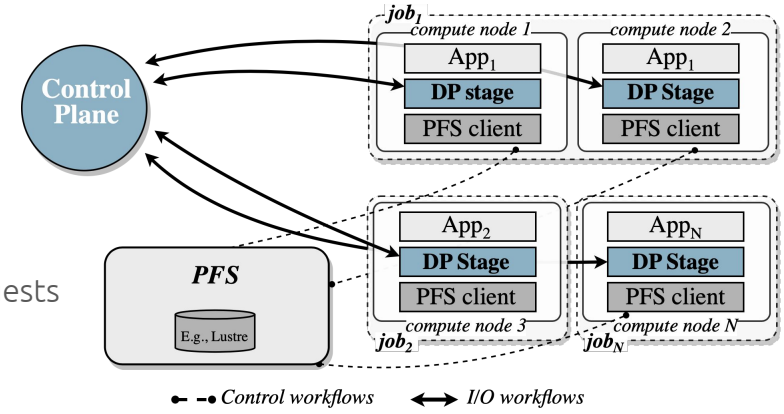

Development Challenges

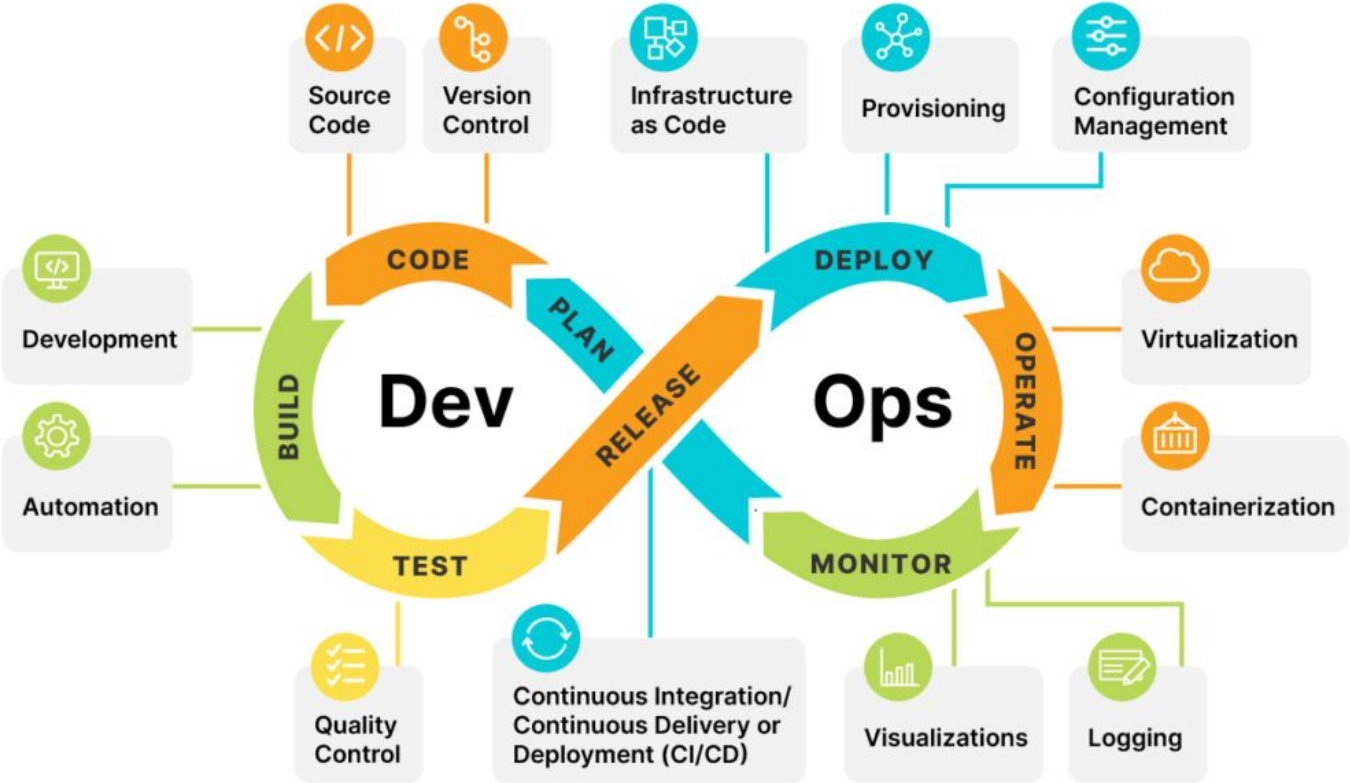

Integration: delivery workflow over testbeds

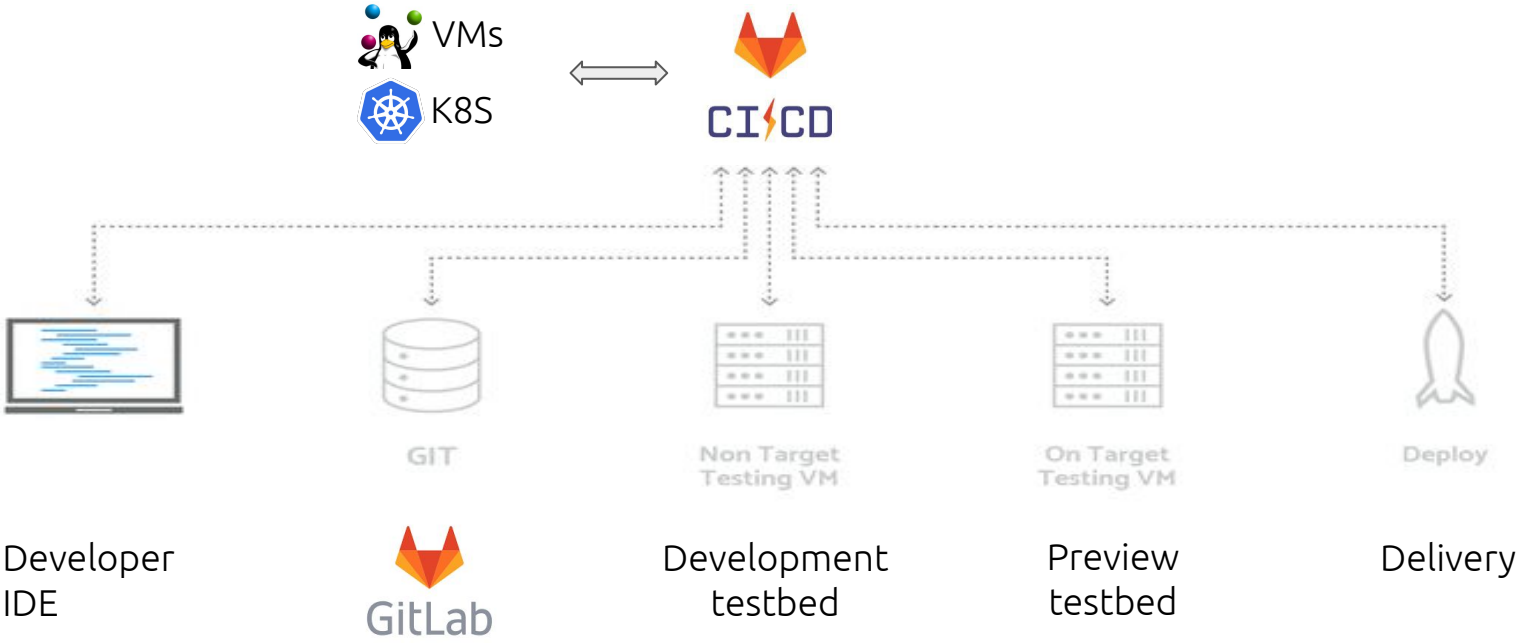

Integration: continuous development

- We use Gitlab cloud platform to create the pipelines to automate all the procedures between issue and code management.

- All the tests are conducted over Gitlab CI pipelines using the provided tools. The output of those tools are summarized in a report accessible from the pipeline job interface.

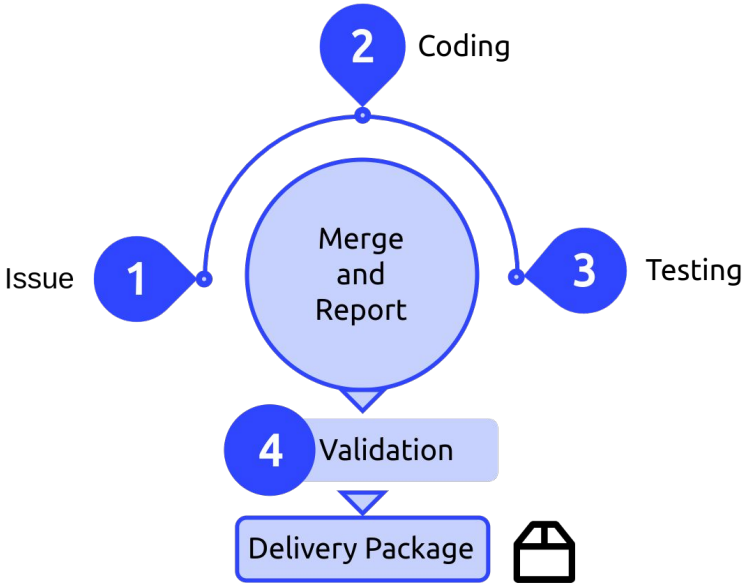

Integration: continuous development

- Create a new project in Gitlab
- Add pipeline from template

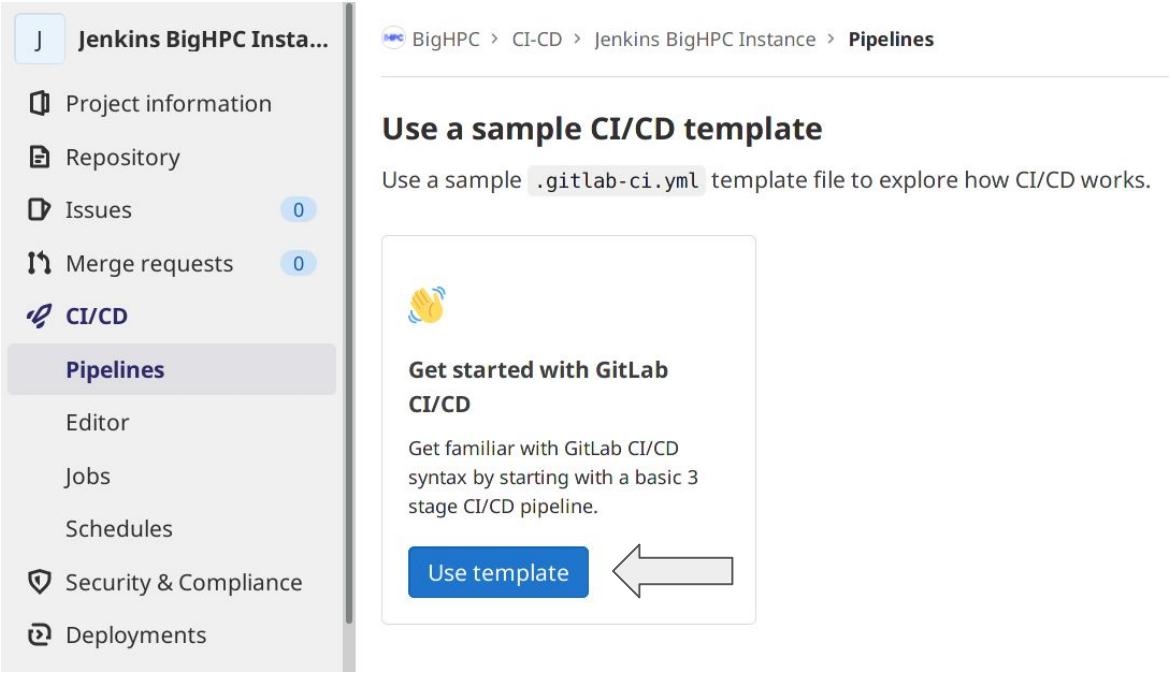

Integration: continuous development

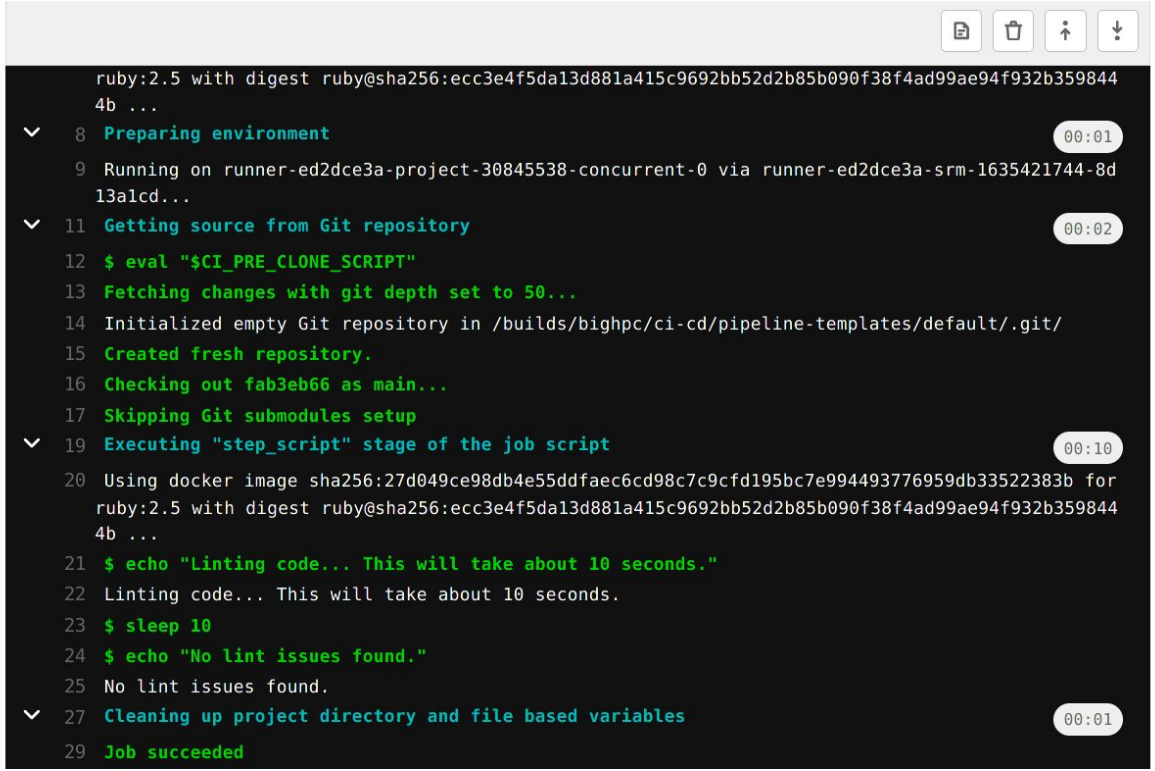

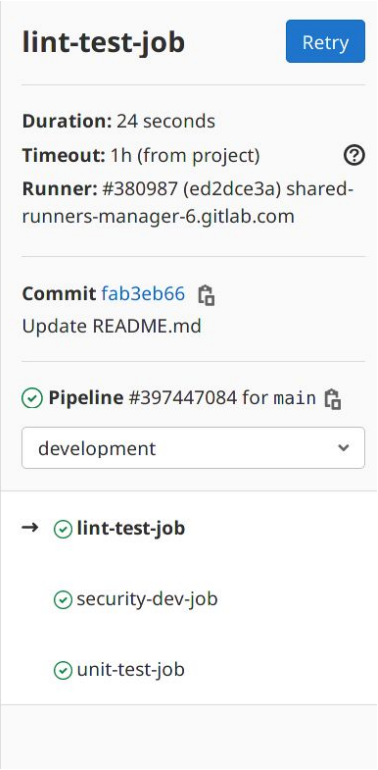

Integration: continuous quality assurance

Involved development and operation domains on the continuous loop:

- Testing, Validation and Delivery
- Deployment and Analysis
- Management and Monitoring

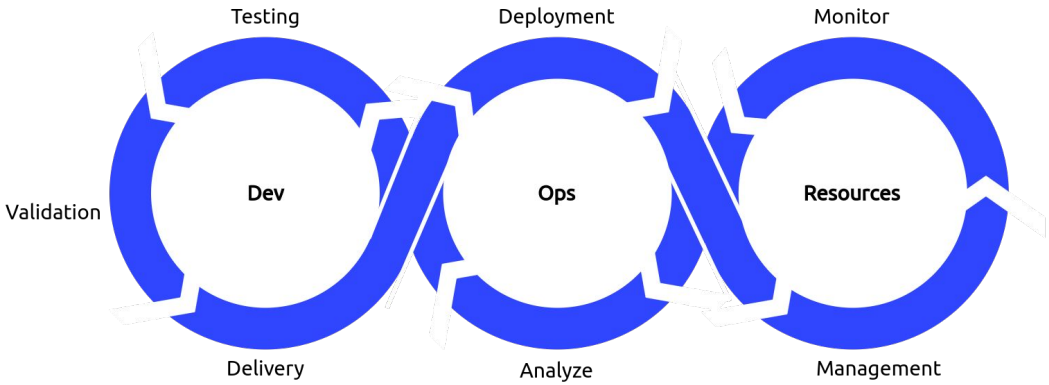

Integration: testing, validation and delivery

Validation is completed in three steps:

- tools running over the pipeline pass the defined tests
- reports are validated by the developers and operators
- merge requests are approved and a new release delivered

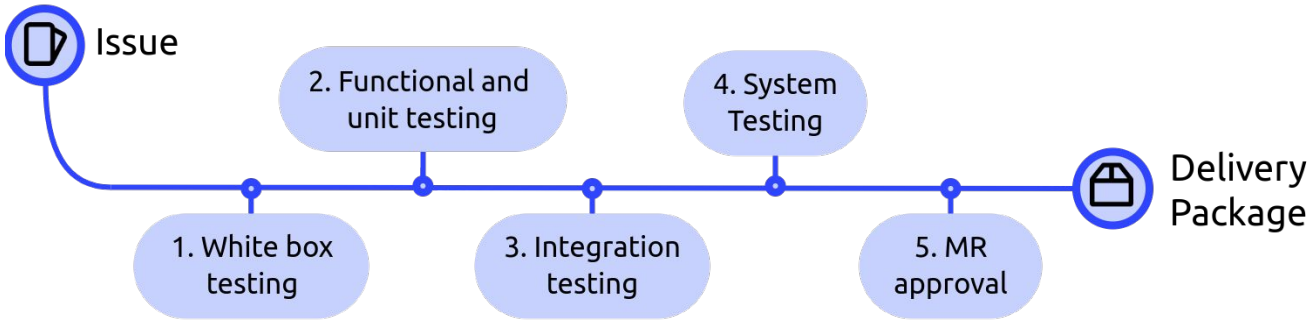

Integration: deployment, analysis and monitor

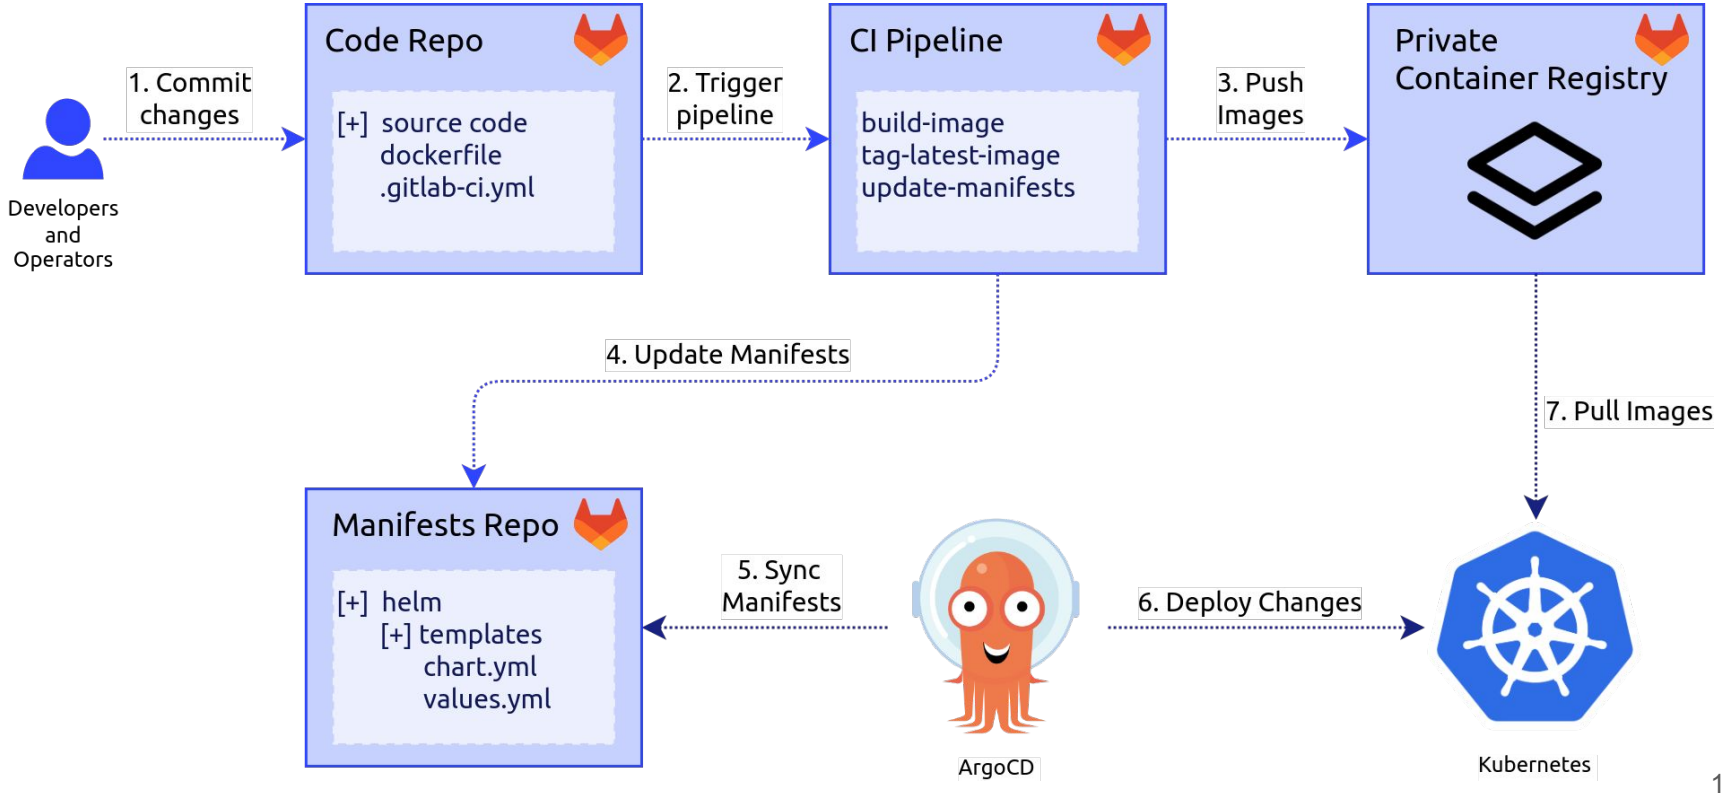

Integration: pilot overview using GitOPS

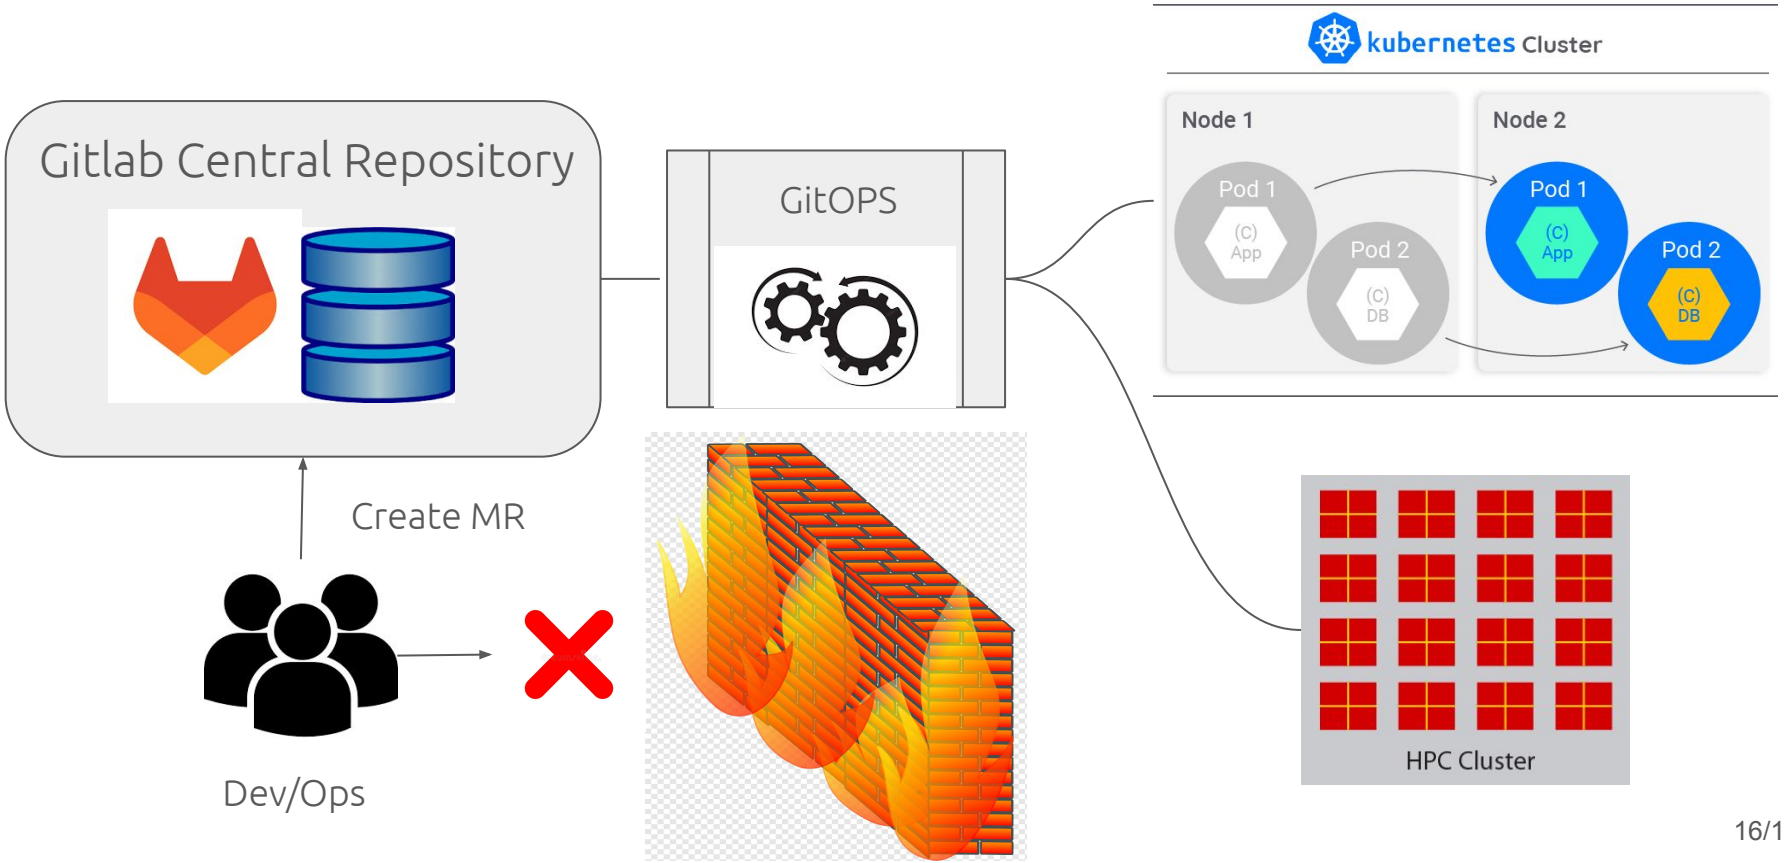

BigHPC platform interfaces

Container Image Upload

File path: *
Path of file to upload [required]

Application name: *
Name of application that the container is built for.

Container name: *
Unique name for container

Library: *
Library information for the container in json format

Container type: *
- Singularity (SIF)
- Docker

CPU instruction set: *
- avx2
- avx512f

Accelerator:
Accelerator that container supports

Upload | Clear fields

Project developed by BigHPC Consortium.

Submit a BigHPC job

Container arguments:
Pass through these arguments to singularity

Name: *
Name of the job

Container image id: *
Container Image Id

Time: *
Time in minutes

Nodes: *
Number of nodes

Processes: *
Number of processes

PPN: *
Processes per node

Partition id: *
Cluster id

Input directory: *
Input directory relative to your root directory in bighpc datastore

Command: *
Execute a command in the container

Launch | Clear fields

Project developed by BigHPC Consortium.

Admin: set BigHPC partition

Partition name: *
Name of the partition

CPU: *
CPU architecture

Site: *
TACC

Memory: *
Memory size in GB

Sockets: *
Sockets number

Accelerator:
Sockets number

Scratch directory: *
Temporary input and output storage path

Login node: *
The FQDN of the login node for the partition

Setup path: *
Path in partition where BigHPC components are installed

Global SDS: *
SDS global controller name and port

Create | Clear fields

Project developed by BigHPC Consortium.

Takeaways

- Unified framework for running parallel and Big Data jobs at HPC centres
- Container-based solution to ease portability and deployment
- Storage QoS policies to ensure fairness and avoiding I/O bottlenecks
- Monitoring and visualization tools to understand cluster's utilization

Questions and Answers time

Thank you for your attention!

More details at: https://bighpc.wavecom.pt/

Contact: info.bighpc@wavecom.pt

BIG HPC HIGH PERFORMANCE COMPUTING

Partners:

INESCTEC

LIP

TEXAS
The University of Texas at Austin

MACC Minho Advanced Computing Center

TACC

wavecom

Funding:

Cofinanciado por:
COMPETE 2020
Lisb@20^20
PORTUGAL 2020
UNIÃO EUROPEIA
Fundo Europeu de Desenvolvimento Regional
FCT Fundação para a Ciência e a Tecnologia
UTAustin Portugal