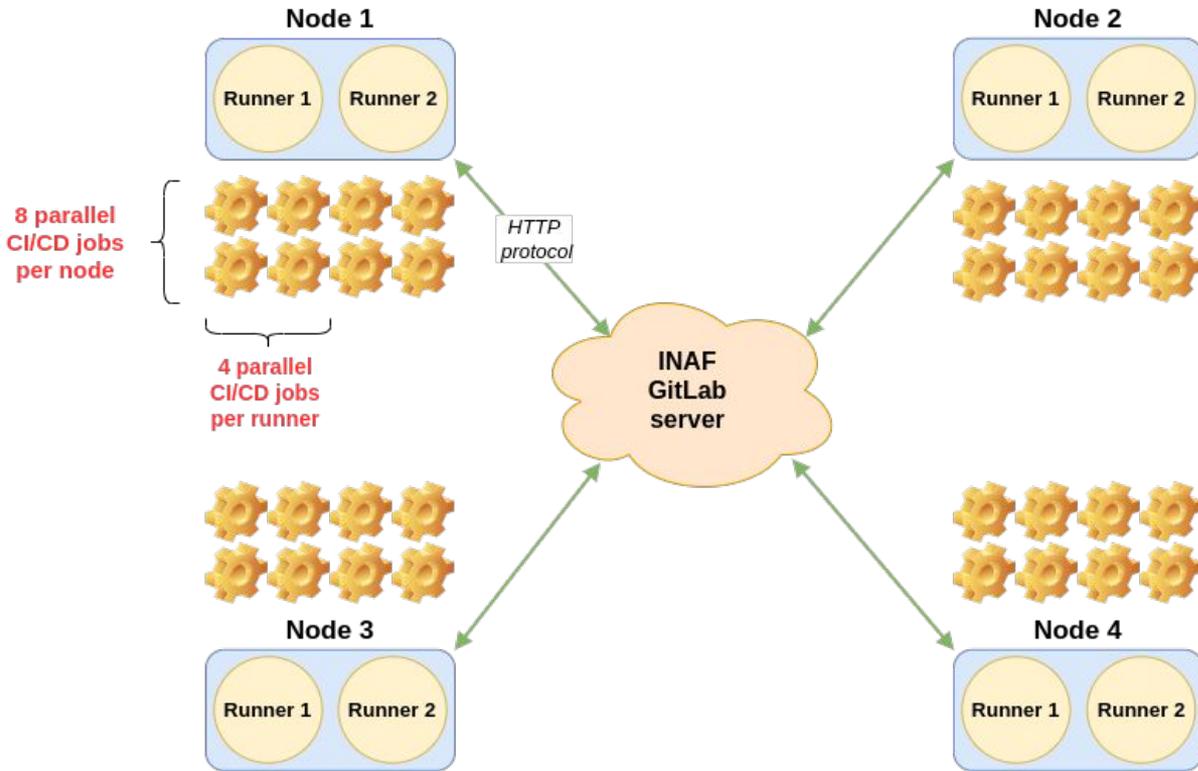


IA2 ongoing developments

Massimo Sponza

Test cluster for GitLab CI/CD runners



Motivation:

increased demand for a GitLab CI/CD support infrastructure from INAF researchers and technologists

Initial test platform:

- 4 physical machines with a Intel(R) Xeon(R) CPU E5645 @ 2.40GHz (1 thread/core, 6 cores/socket and 2 sockets) and 96 GB of RAM, kindly provided by the OATs Planck group (thanks to A. Zacchei)

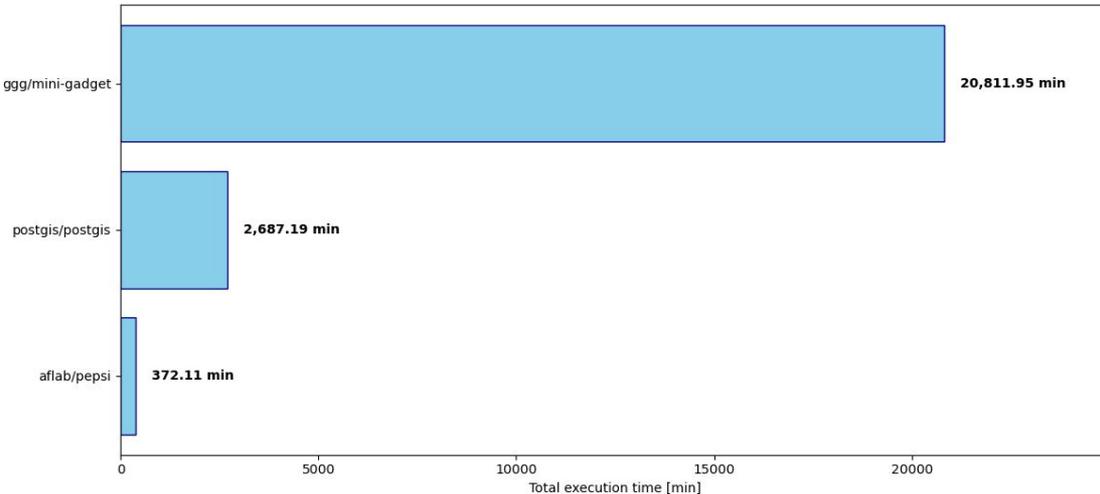
Production cluster:

- now virtualized on new hardware: up to 48 vCores and up to 256 GB RAM

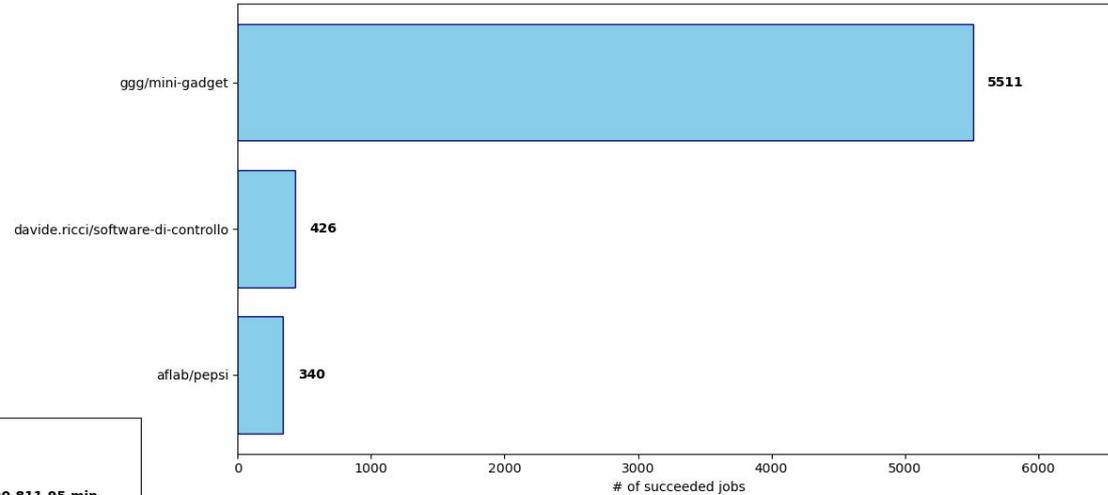
Test cluster for GitLab CI/CD runners: some stats

- **Total number of jobs:** ~11.5k
- **Jobs success rate:** ~70%

Top 3 projects in terms of total execution time



Top 3 projects in terms of number of succeeded jobs

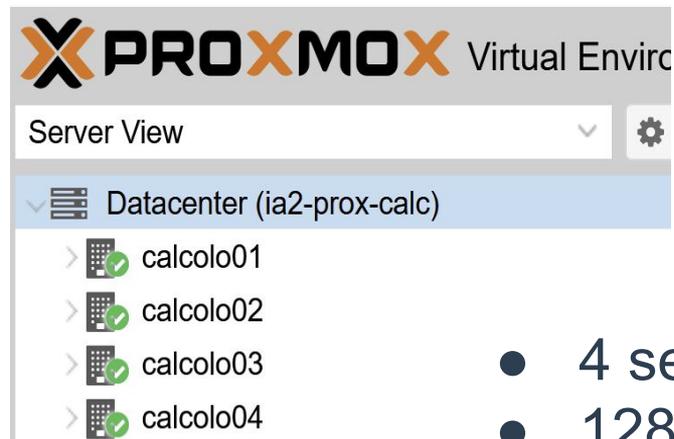


NOTE: the stats showed in these charts consider only real projects (no tests, schools or courses)

Rosetta worker nodes

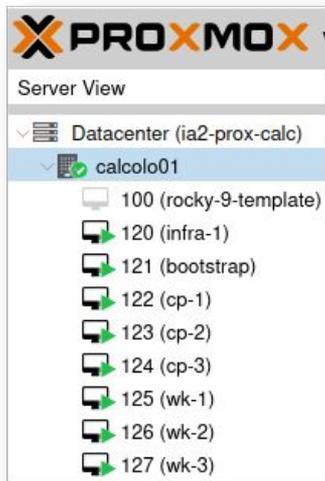
- need to perform small computing tasks on users data
- GitLab runners are not designed to provide access to user data
- Solution: integrate some of the nodes as Rosetta worker nodes (thanks to Stefano Alberto Russo)
- User get personal home directory mounted as volume in a container
- Singularity
 - it works for user data
 - multiple batch sessions per node
 - limitations: only one graphical session per node

New Hardware and Proxmox



- 4 servers
- 128 vCores per server
- 2 TB RAM per server
- CEPH Hyper Converged Storage, 6 SSDs per server, 24 TB disk space available
- GPU NVIDIA Ada L40S 48GB GDDR6
- Proxmox VE, alternative to VMware
- VMware will be dismissed in INAF

OKD - K8s integration on Proxmox VE



The screenshot displays the OKD web console interface. The top navigation bar shows the 'okd' logo and the user 'kube-admin'. The main content area is titled 'Overview' and provides a high-level view of the cluster. It includes sections for 'Getting started resources', 'Status', 'Cluster inventory', and 'Cluster utilization'. The 'Status' section shows the cluster is 'OK' with various components like 'Control Plane', 'Operators', and 'Dynamic Playbooks' in a healthy state. The 'Cluster utilization' section features a graph showing resource usage for CPU, Memory, Filesystem, and Network transfer over time. The 'Activity' section on the right lists recent events, such as 'Server has stopped listening' and 'All pending requests processed'.

Overview

You are logged in as a temporary administrative user. Update the [cluster OAuth configuration](#) to allow others to log in.

Getting started resources

- [Set up your cluster](#)
Finish setting up your cluster with recommended configurations.
 - [Take console tour](#)
 - [Add identity providers](#)
 - [Configure alert receivers](#)[View all steps in documentation](#)
- [Build with guided documentation](#)
Follow guided documentation to build applications and familiarize yourself with key features.
 - [Enable the Developer Perspective](#)
 - [Impersonating the system:admin user](#)[View all quick starts](#)
- [Explore new features and capabilities](#)
OpenShift AI
Build, deploy, and manage AI-enabled applications.
 - [Trusted Software Supply Chain](#)
Automatically validate integrity, source and risks, release safety.
 - [French and Spanish now available](#)
Console language options now include French and Spanish.[See what's new in OpenShift 4.21](#)

Details [View settings](#)

Cluster API address
<https://api.okd.ia2.inf.r.t.6443>

Cluster ID
3454ebbc-51a2-45b9-9ae7-cdf2e6e38d9

Infrastructure provider
None

OpenShift version
4.21.0-okd-scos.7

Update channel
stable-scos-4

Status [View alerts](#)

Cluster Control Plane Operators Dynamic Playbooks Insights Disabled

AlertManagerReceiversNotConfigured [Configure](#)
Mar 5, 2025, 4:29 PM
Alerts are not configured to be sent to a notification system, meaning that you may not be notified in a timely fashion when important failures occur. Check the OpenShift documentation to learn how to configure notifications with Alertmanager.

KubeDaemonSetMisScheduled [View details](#)
Mar 5, 2025, 4:28 PM
3 Pods of DaemonSet operator/ingress-canning/ingress-canary are running where they are not supposed to run.

KubeDaemonSetMisScheduled [View details](#)
Mar 5, 2025, 4:28 PM
3 Pods of DaemonSet operator/insights/insights-runtime-extractor are running where they are not supposed to run.

Cluster inventory [View settings](#)

[6 Nodes](#)
[301 Pods](#) 23 1
[0 StorageClasses](#)
[0 PersistentVolumeClaims](#)

Cluster utilization [View settings](#)

Resource Usage

CPU
22.7% available of 24 1.2%
10
5

Memory
6724 GiB available of 93.66 GiB 26.42 GiB
40 GiB
20 GiB

Filesystem
624 GiB available of 718.7 GiB 94.7 GiB
100 GiB
50 GiB

Network transfer 17 Mbps in 17 Mbps out
20 Mbps

Pod count 200
200

Activity

Ongoing

There are no ongoing activities.

Recent events [Pause](#)

- Server has stopped listening
- All pending requests processed
- All pre-shutdown hooks have been finished
- Received signal to terminate, becoming unready, but keeping serving
- All pre-shutdown hooks have been finished
- All pending requests processed
- Server has stopped listening
- Received signal to terminate, becoming unready, but keeping serving
- All pending requests processed

[View all events](#)