

Recent Developments in JINR Cloud Services

N. Balashov, A. Baranov, N. Kutovskiy, A. Makhalkin, Ye. Mazhitova, R. Semenov

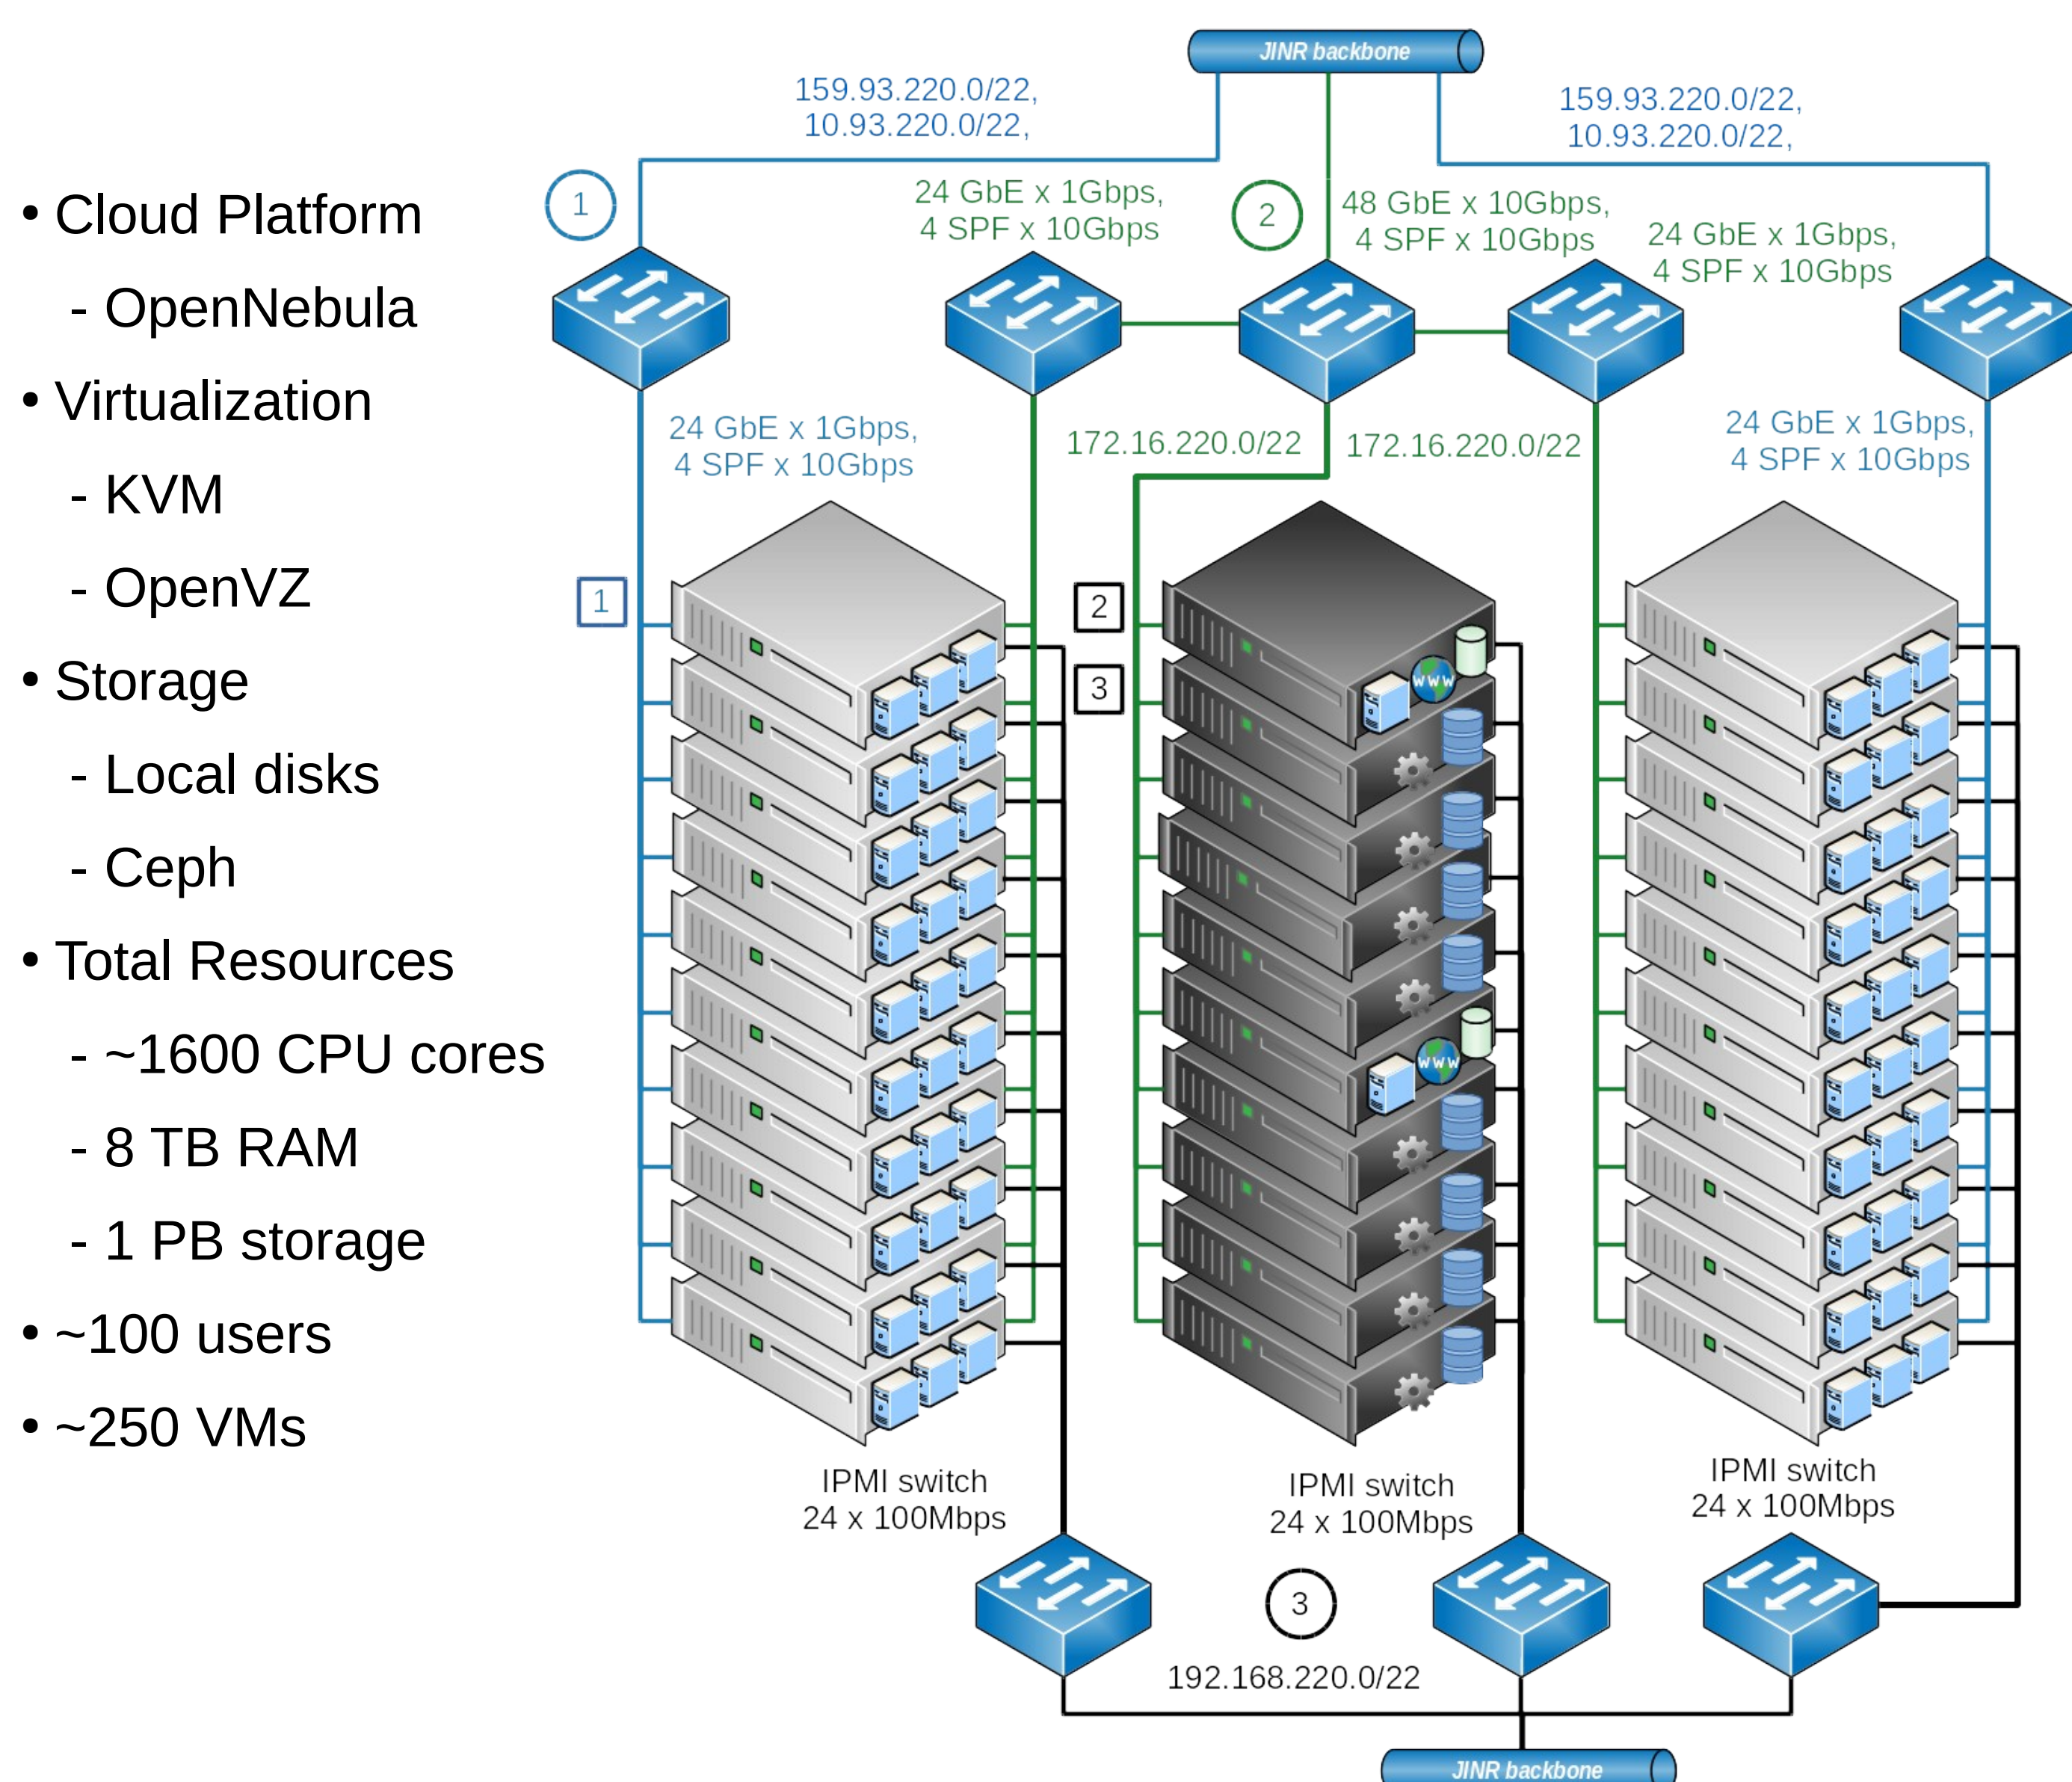
Laboratory of Information Technologies, JINR



INTRODUCTION

The JINR cloud service, based on the Infrastructure-as-a-Service (IaaS) model, is used as a universal computing resource that supports individual scientists and provides a variety of common services, ranging from simple web-sites to complex multiuser computational systems. In this paper we give an overview of the recent changes in the JINR cloud infrastructure and present new common cloud services: `mon-service.jinr.ru` and `jupyter.jinr.ru`

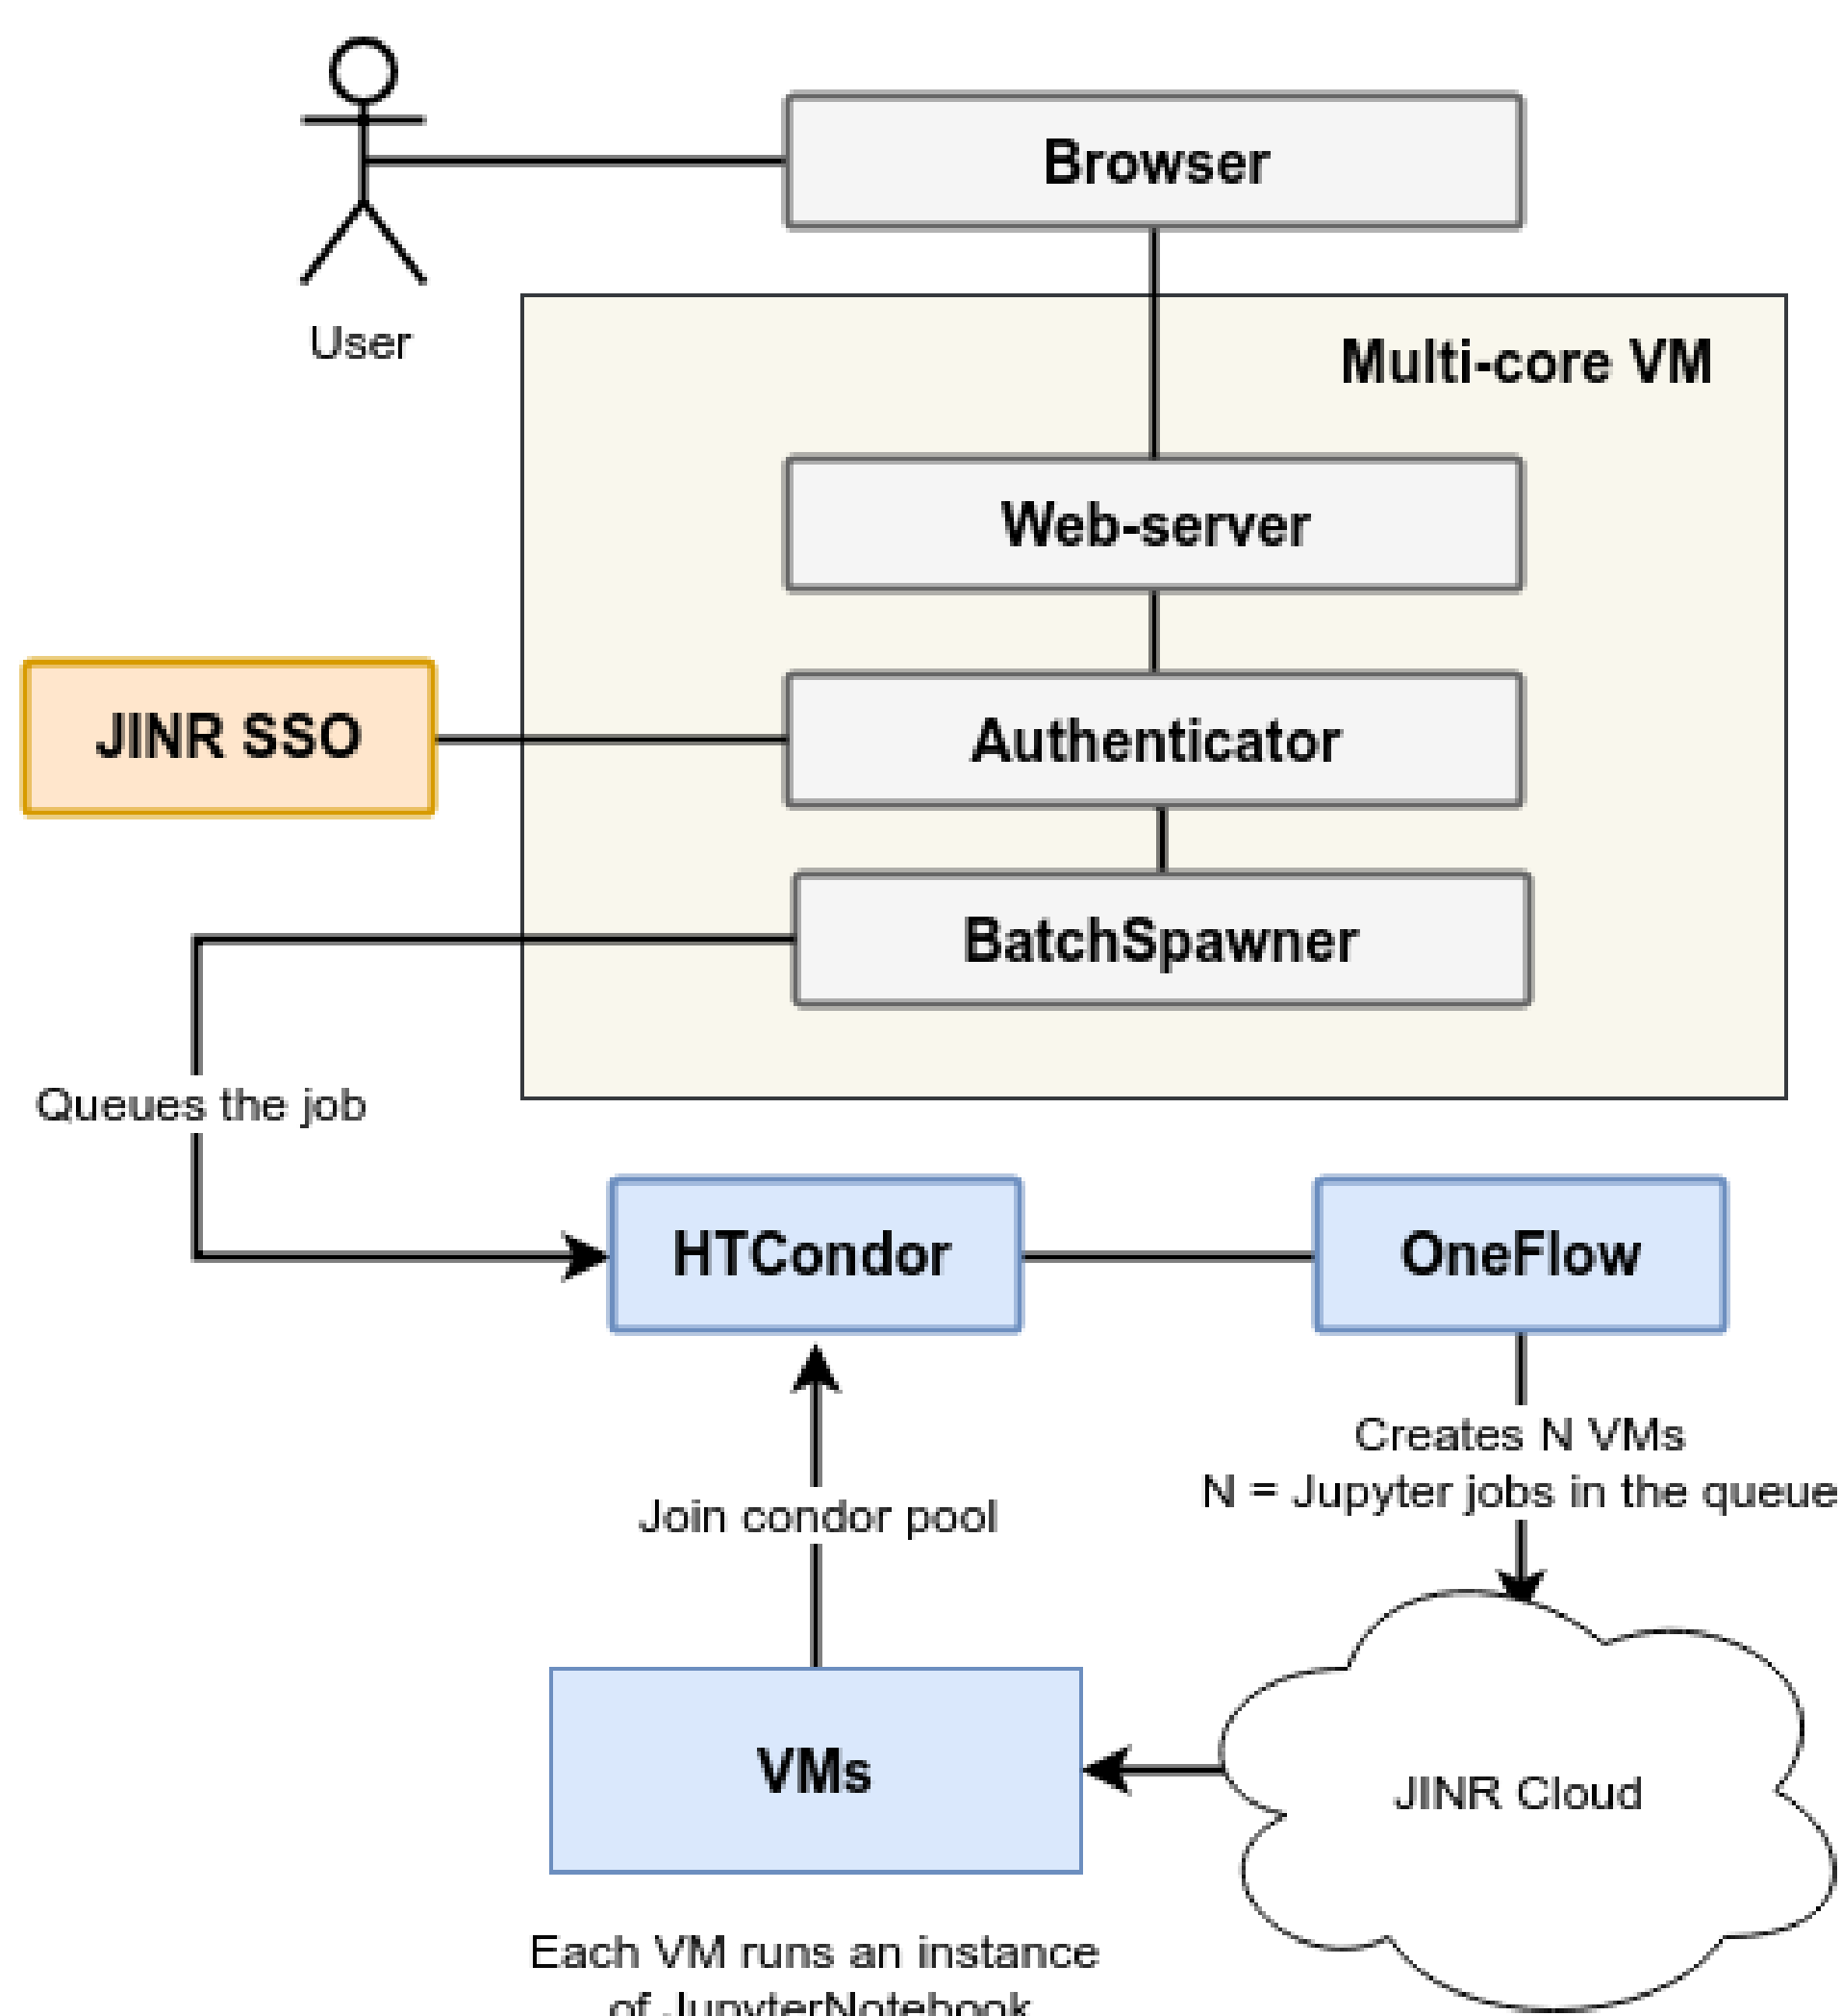
JINR Cloud Overview



The cloud provides virtual machines (VMs) and containers for personal usage and as a hosting platform for multiuser services. The resources are mostly used by software developers, system administrators and physicists from JINR and its partner organizations.

JupyterHub Service

JupyterHub is a multi-user service providing Jupyter Notebooks – documents that contain live code, equations, visualizations and narrative text. Uses include data analysis and educational purposes.



Currently the JupyterHub cluster has a static configuration with a small set of worker-nodes connected with the NOVA environment set up. This service is operating in testing mode and is used only for a small subset of data analysis and educational tasks by the NOVA experiment team. Due to the fixed number of computational nodes and a highly specific use-cases it is often underutilized and, therefore, is wasting computing resources. To deal with this problem we are implementing a dynamic (“elastic”) system, the scheme of which is shown above.

The main components of the new system that make it “elastic” are:

- The BatchSpawner which enables JupyterHub to use HTCondor as a computational backend.
- JINR Single Sign-On (SSO) is used to authenticate users – no need to create special accounts.
- VMs in the cloud are used as HTCondor worker-nodes.
- OneFlow service handles creation/deletion of the VMs according to demand.

Introduction of this automatic system will help better align the amount of computing resources with the actual demand, thus reducing the underutilization of cloud resources by the JupyterHub service.

Monitoring Service

A new monitoring service was launched recently – `mon-service.jinr.ru`. The service is based on the Grafana platform which provides powerful tools for visualizing the monitoring data. The supported data back-ends include most of the commonly used storage systems: Graphite, InfluxDB, Prometheus and others. An example dashboard is shown on the picture below.



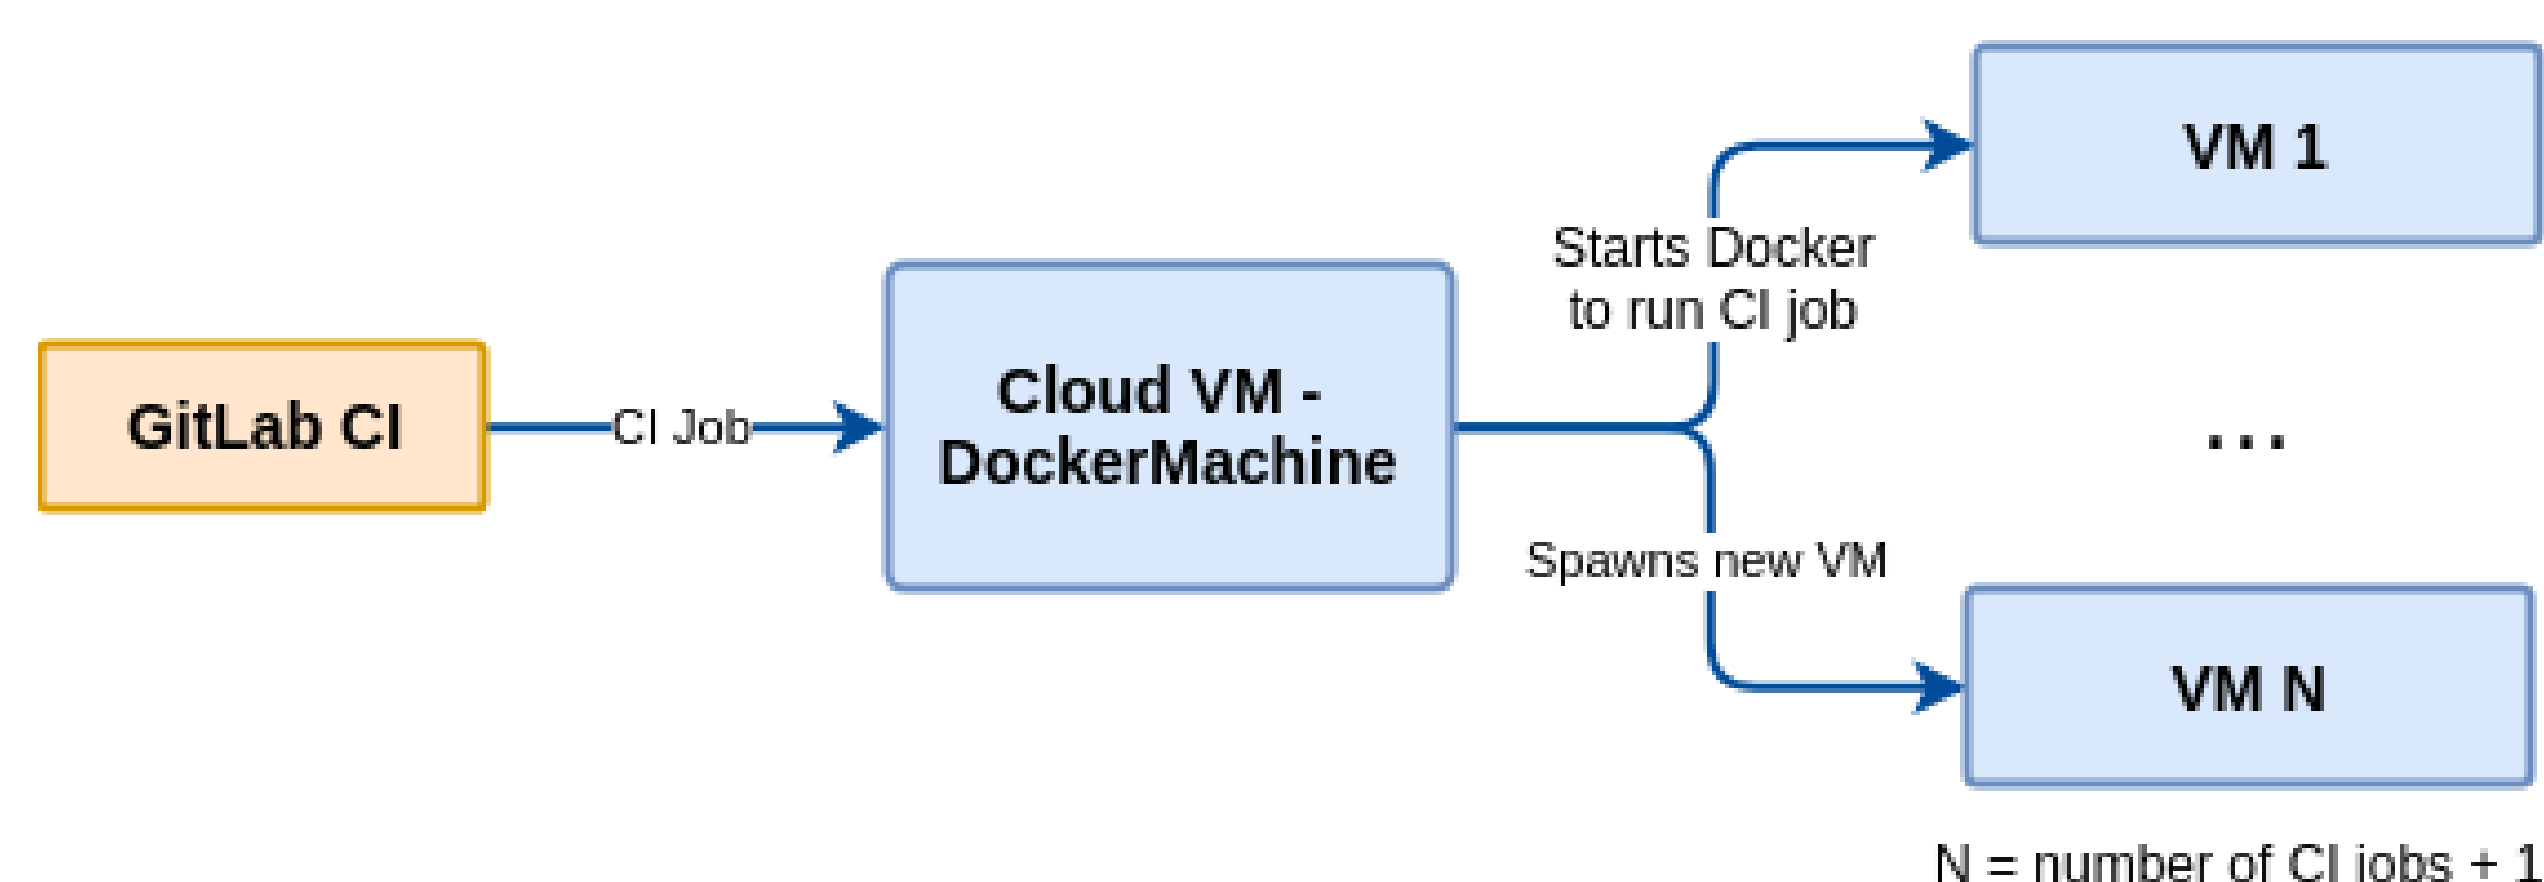
The service is connected to the JINR SSO for the authentication and can be used by the JINR system administrators. It is currently used to monitor the following systems:

- The `git.jinr.ru` service.
- The Remote Operations Center (ROC) Dubna of the NOVA experiment.
- The JINR Cloud.

Scalable GitLab Runner

The `git.jinr.ru` service, built on the GitLab platform, has a built-in support for Continuous Integration (CI), which requires so called GitLab Runners to be attached to the service. The Runners are the computing resources that execute the actual CI jobs.

To provide our users with computing resources for their CI jobs we deployed an «elastic» Runner and its operation scheme is shown below.



This runner is shared, meaning it can be used by any project at `git.jinr.ru`, and it scales automatically to keep up with the demand:

- CI jobs run in Docker containers, which provide the execution environment.
- The JINR cloud VMs host the Docker containers.
- The special DockerMachine service handles on-demand creation/deletion of the VMs in the cloud.