



Laniakea - Update 2021

Tangaro M.A, Donvito G., Antonacci M., Chiara M.,
Mandreoli P., Pesole G., Zambelli F.

Galaxy Community Conference (Virtual edition)
28 June - 10 July 2021



Laniakea

LANIAKEA is a cloud based Galaxy instance provider.

<https://laniakea-elixir-it.github.io/>

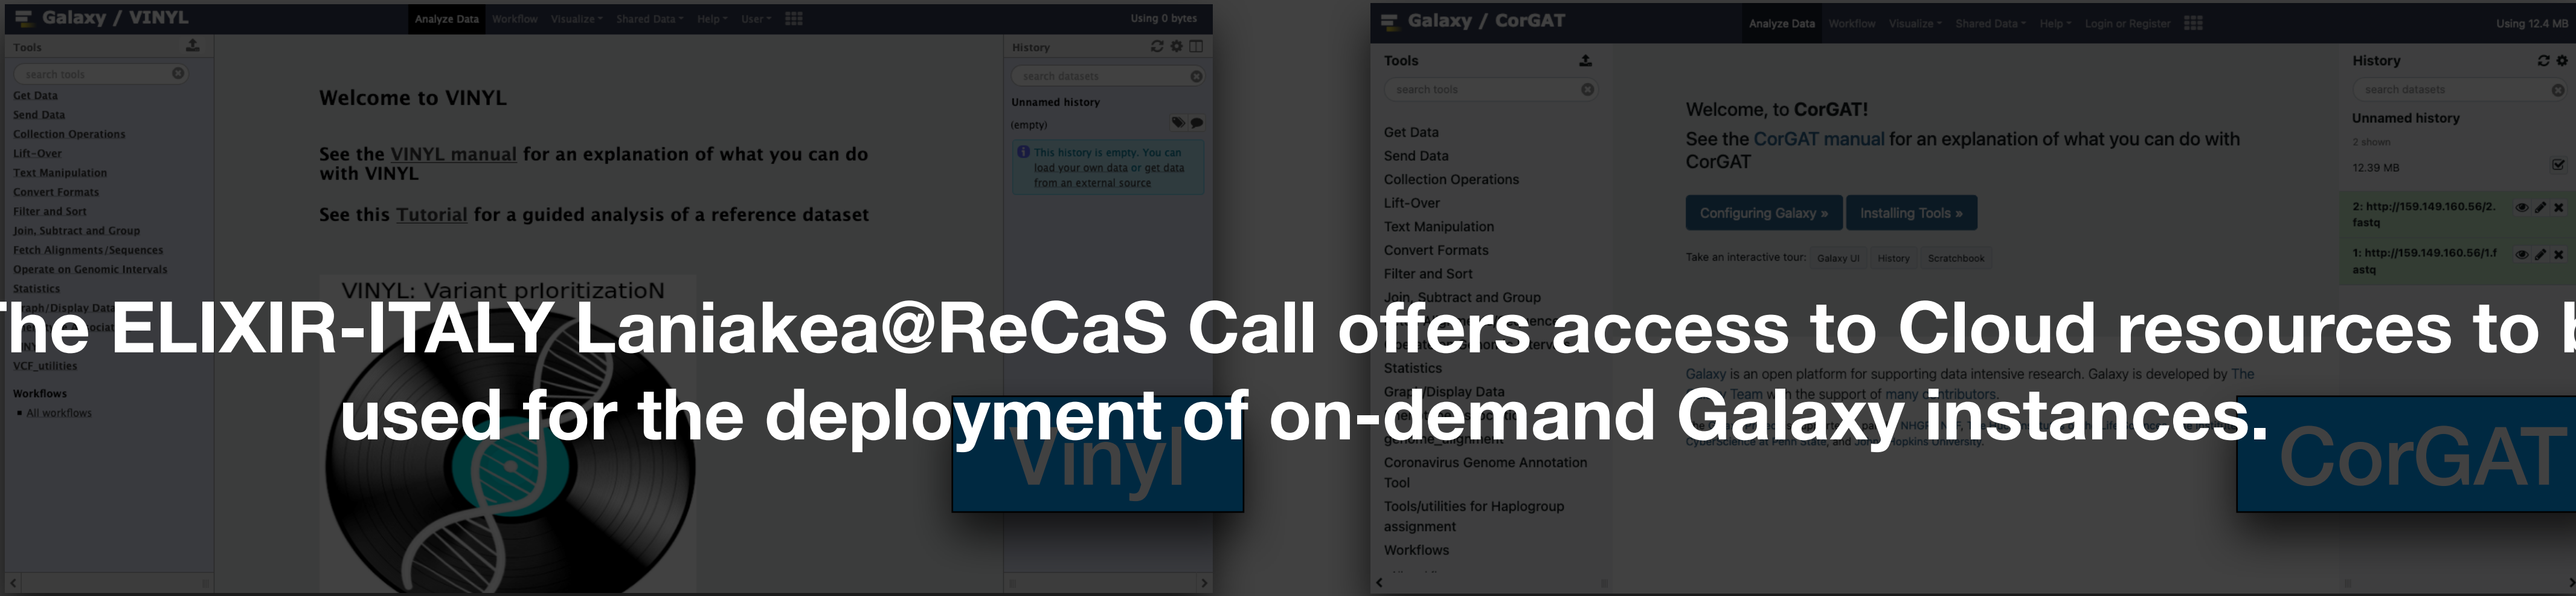
It relies on INDIGO-DataCloud PaaS layer, GalaxyProject tools, e.g. the Official Galaxy Docker Container and Ephemervis and common Open Source tools, e.g. HashiCorp Vault, LUKS and SLURM.

Recommended for scenarios where users need full administrative control over a private Galaxy instance.



One year of Laniakea@ReCaS

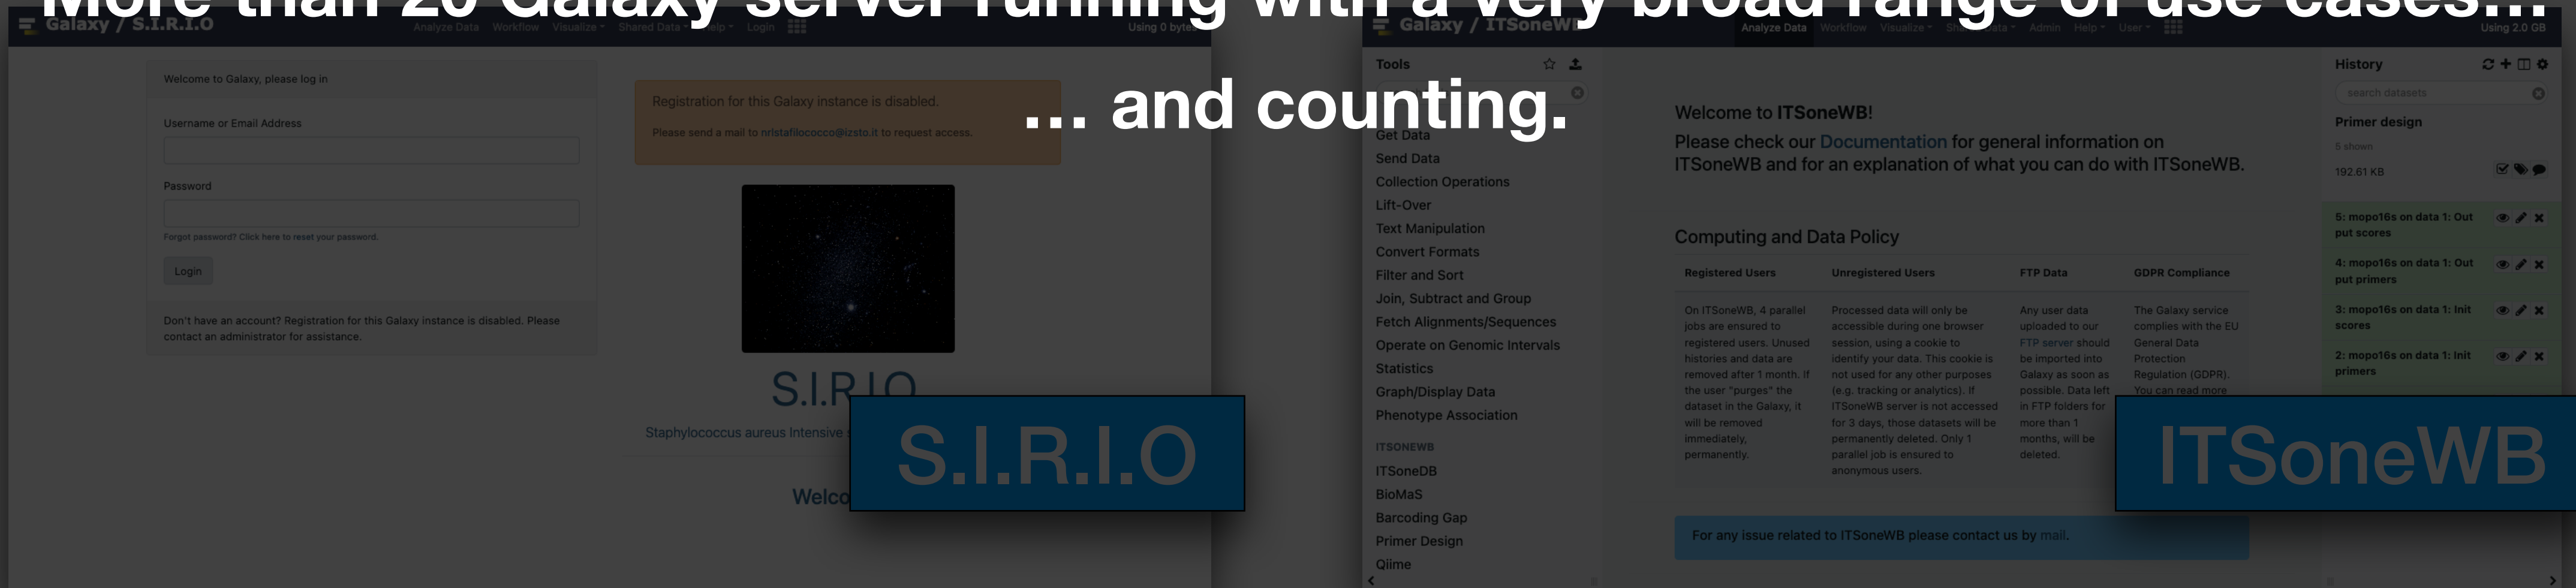
The ELIXIR-ITALY Laniakea@ReCaS Call offers access to Cloud resources to be used for the deployment of on-demand Galaxy instances.



The image shows two screenshots of Galaxy web interfaces. The left screenshot is for 'Galaxy / VINYL' and the right is for 'Galaxy / CorGAT'. Both show a 'Welcome' message and a list of tools on the left sidebar. The VINYL instance has a 'History' panel showing an empty history. The CorGAT instance has a 'History' panel showing two entries: '2: http://159.149.160.56/2.fastq' and '1: http://159.149.160.56/1.1.fastq'. A blue box with the text 'vinyl' is overlaid on the bottom right of the VINYL screenshot, and a blue box with the text 'CorGAT' is overlaid on the bottom right of the CorGAT screenshot.

More than 20 Galaxy server running with a very broad range of use cases...

... and counting.

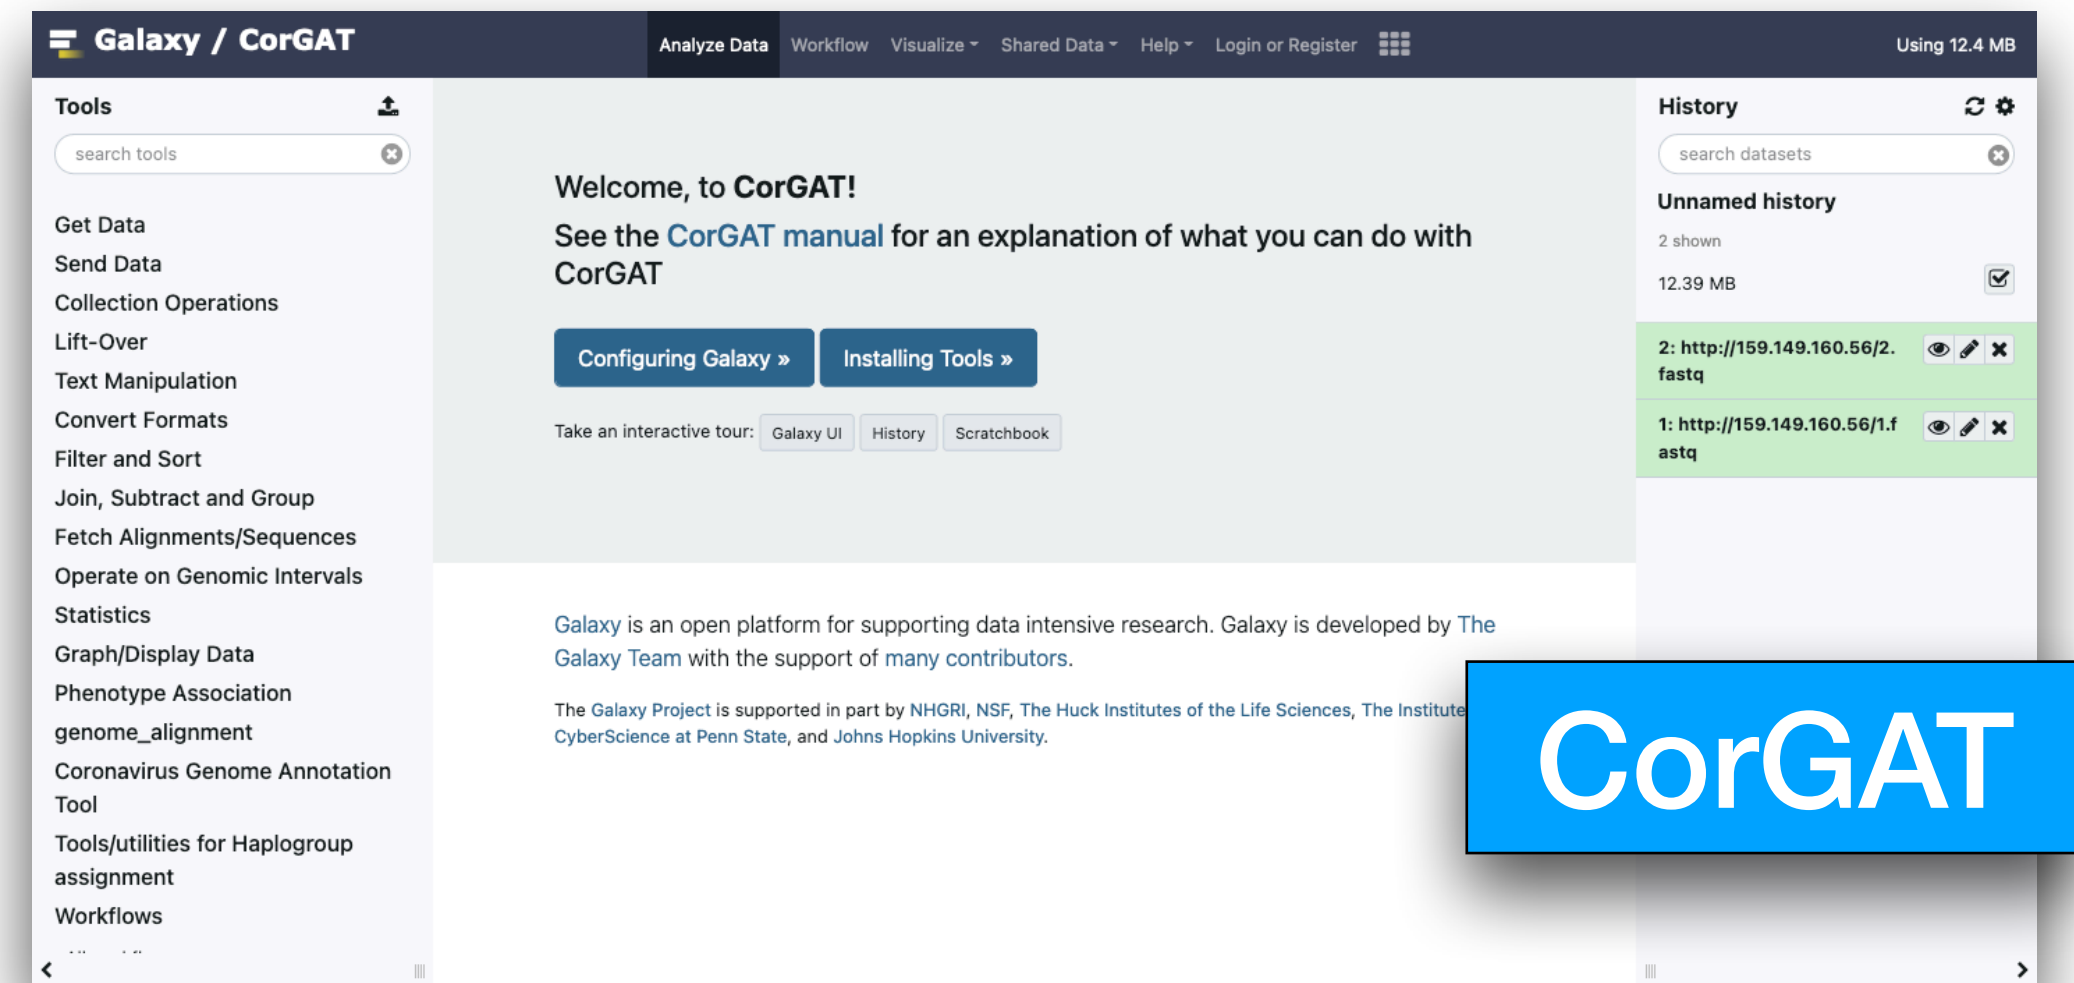


The image shows two screenshots of Galaxy web interfaces. The left screenshot is for 'Galaxy / S.I.R.I.O' and the right is for 'Galaxy / ITSoneWB'. The S.I.R.I.O instance shows a login form with fields for 'Username or Email Address' and 'Password', and a 'Login' button. A message above the form says 'Registration for this Galaxy instance is disabled. Please send a mail to rristaffiloccoco@izsto.it to request access.' The ITSoneWB instance shows a 'Welcome to ITSoneWB!' message and a 'Primer design' tool in the history panel. A blue box with the text 'S.I.R.I.O' is overlaid on the bottom right of the S.I.R.I.O screenshot, and a blue box with the text 'ITSoneWB' is overlaid on the bottom right of the ITSoneWB screenshot.

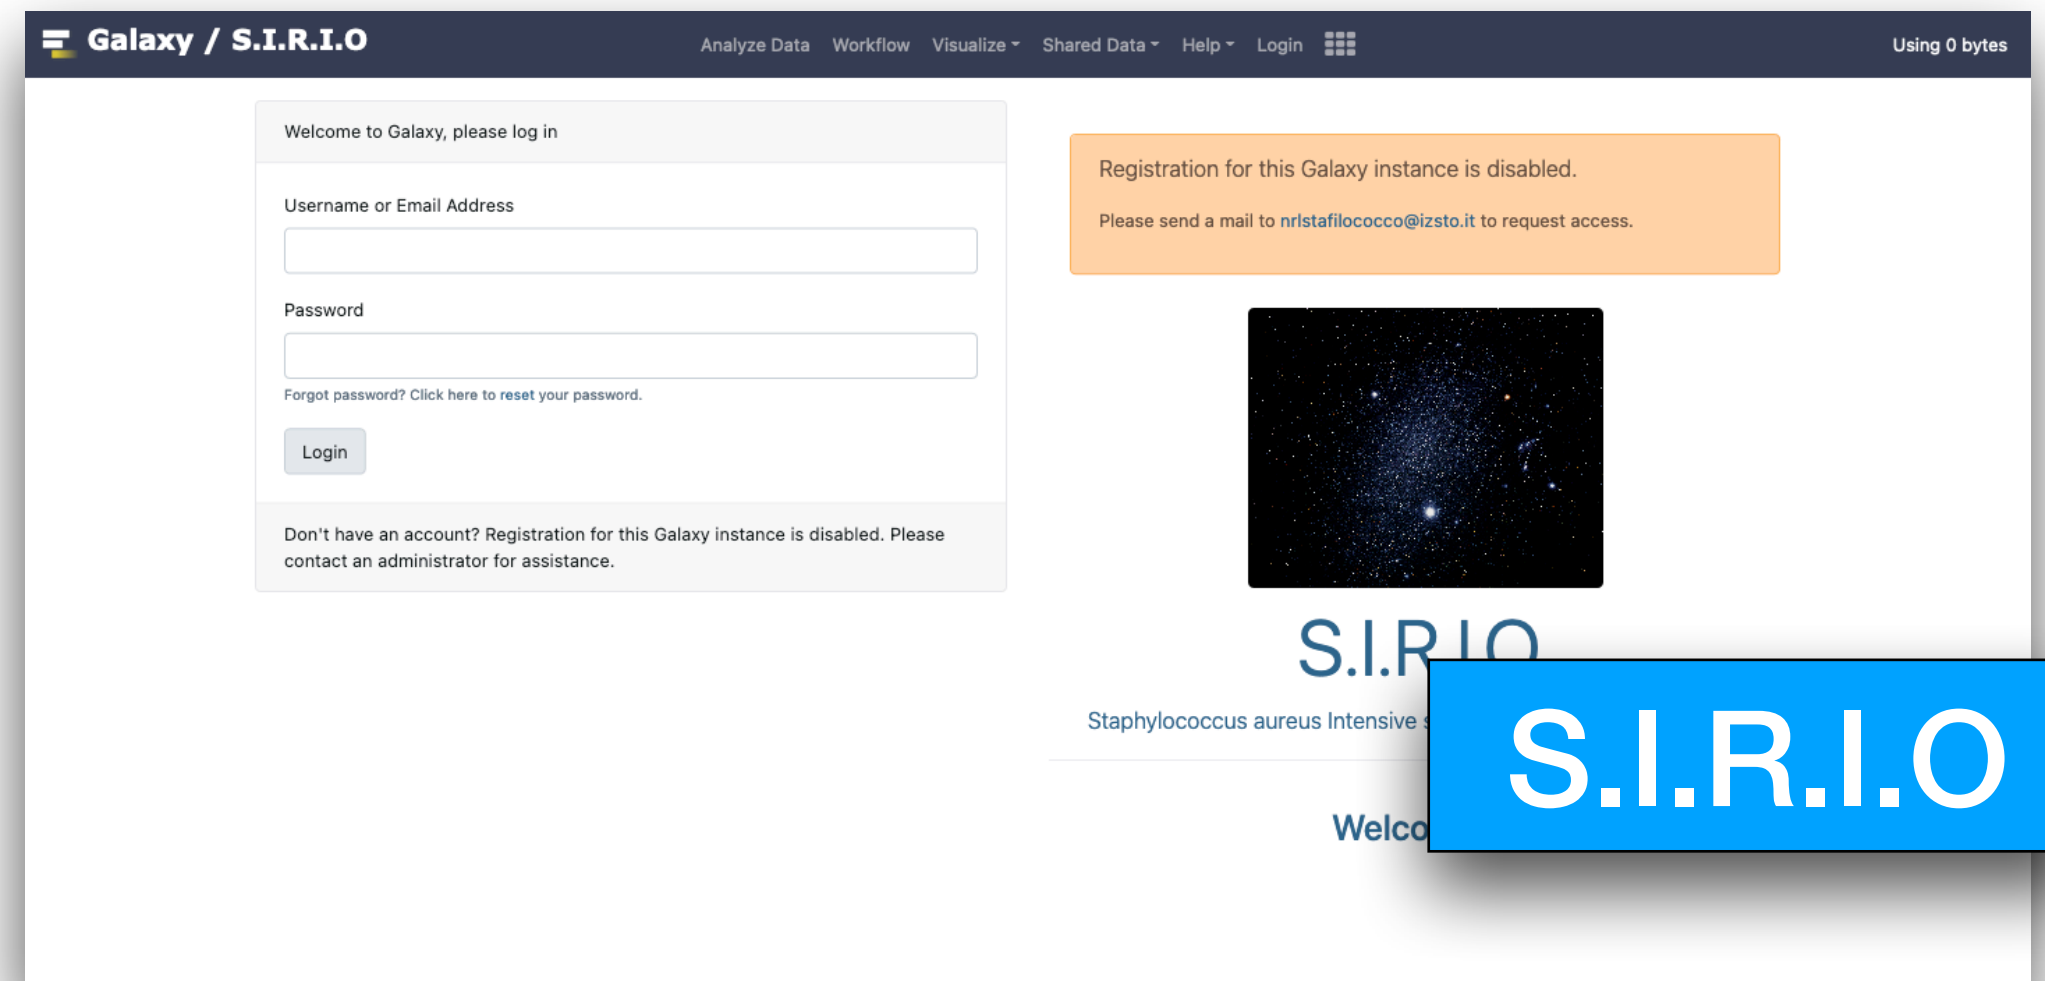
One year of Laniakea@ReCaS



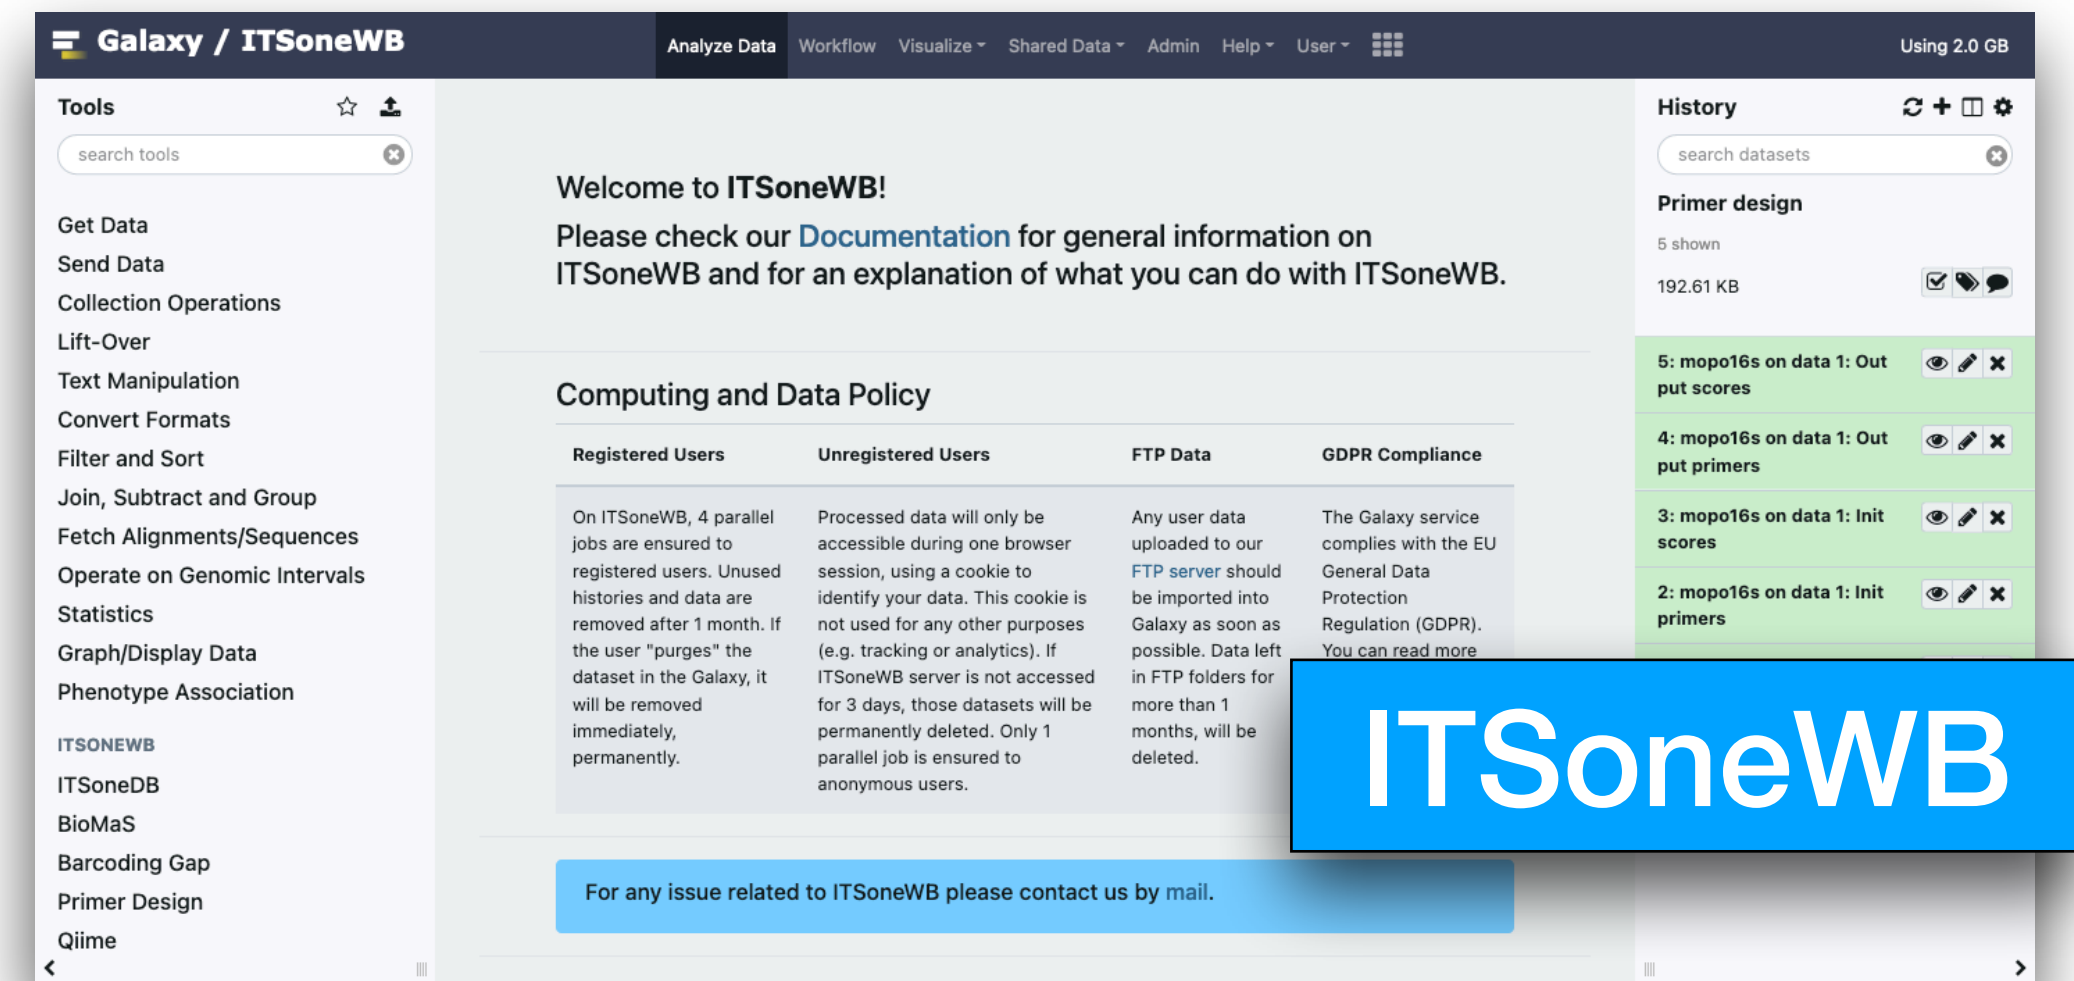
The screenshot shows the Galaxy VINYL interface. The main content area displays a welcome message: "Welcome to VINYL". Below this, it says "See the [VINYL manual](#) for an explanation of what you can do with VINYL" and "See this [Tutorial](#) for a guided analysis of a reference dataset". A central graphic features a vinyl record with a DNA double helix overlaid on it, titled "VINYL: Variant prioritization". A blue box with the text "Vinyl" is overlaid on the bottom right of the screenshot.



The screenshot shows the Galaxy CorGAT interface. The main content area displays a welcome message: "Welcome, to CorGAT!". Below this, it says "See the [CorGAT manual](#) for an explanation of what you can do with CorGAT". There are two buttons: "Configuring Galaxy" and "Installing Tools". Below these, it says "Take an interactive tour: [Galaxy UI](#) [History](#) [Scratchbook](#)". A blue box with the text "CorGAT" is overlaid on the bottom right of the screenshot.



The screenshot shows the Galaxy S.I.R.I.O interface. The main content area displays a login form with fields for "Username or Email Address" and "Password", and a "Login" button. A message box states: "Registration for this Galaxy instance is disabled. Please send a mail to nristafilococco@izsto.it to request access." Below this is a starry space image and the text "S.I.R.I.O" and "Staphylococcus aureus Intensive". A blue box with the text "S.I.R.I.O" is overlaid on the bottom right of the screenshot.



The screenshot shows the Galaxy ITSoneWB interface. The main content area displays a welcome message: "Welcome to ITSoneWB!". Below this, it says "Please check our [Documentation](#) for general information on ITSoneWB and for an explanation of what you can do with ITSoneWB." There is a section titled "Computing and Data Policy" with a table of policies. A blue box with the text "ITSoneWB" is overlaid on the bottom right of the screenshot.

Registered Users	Unregistered Users	FTP Data	GDPR Compliance
On ITSoneWB, 4 parallel jobs are ensured to registered users. Unused histories and data are removed after 1 month. If the user "purges" the dataset in the Galaxy, it will be removed immediately, permanently.	Processed data will only be accessible during one browser session, using a cookie to identify your data. This cookie is not used for any other purposes (e.g. tracking or analytics). If ITSoneWB server is not accessed for 3 days, those datasets will be permanently deleted. Only 1 parallel job is ensured to anonymous users.	Any user data uploaded to our FTP server should be imported into Galaxy as soon as possible. Data left in FTP folders for more than 1 months, will be deleted.	The Galaxy service complies with the EU General Data Protection Regulation (GDPR). You can read more

One year of Laniakea@ReCaS

- Rapid development of bioinformatics tools,
- Efficient delivery of training activities,
- Provision of public bioinformatics services in different settings, including food safety and clinical research.

Laniakea@ReCaS: first year of activity of a Laniakea-based Galaxy “on-demand” service

Pietro Mandreoli
Poster
July 6, 12.00 - 13.00



Working side by side with Researchers and Clinicians, for Galaxy maintenance and tools developments.

New requirements emerged from the user community and cloud providers, driving the development of the next Laniakea release.

Laniakea next release

Platform level

- Python 2 -> Python 3 transition.
- Ansible and package update...

Service level

- New Service Level Agreement tool and Users Quota management.

Development level

- Simplify Galaxy deployment strategy in order to reduce the maintenance effort.

Users level

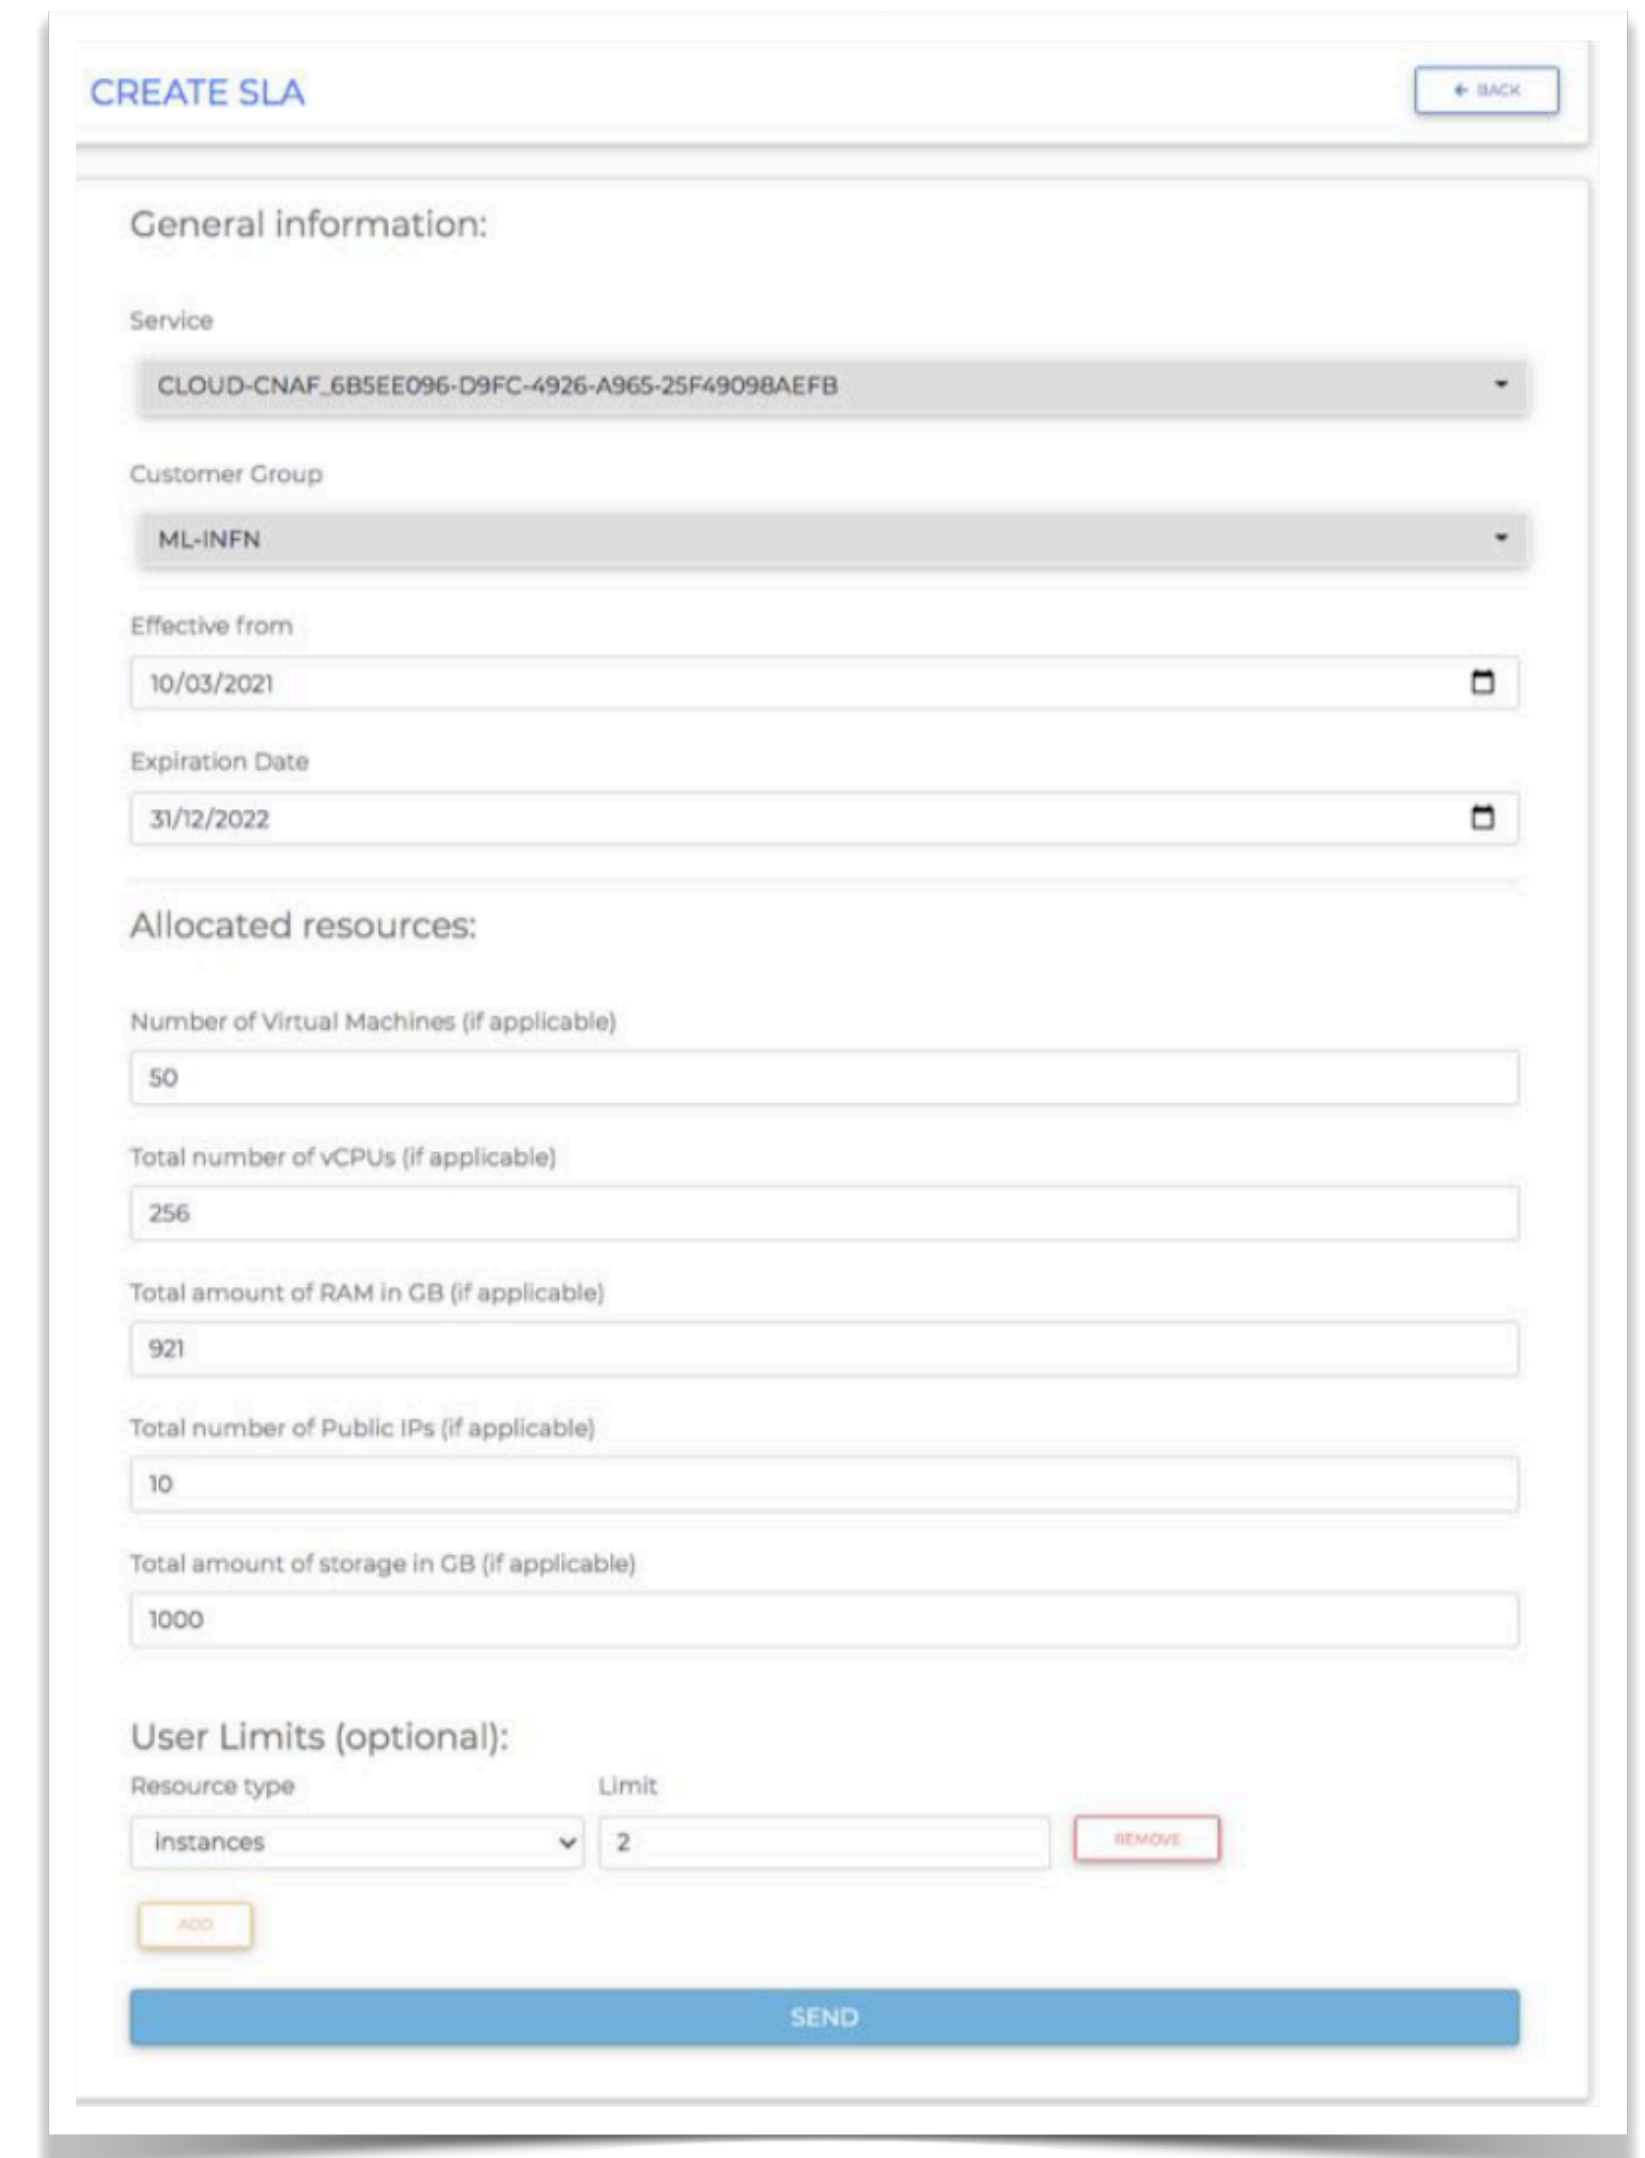
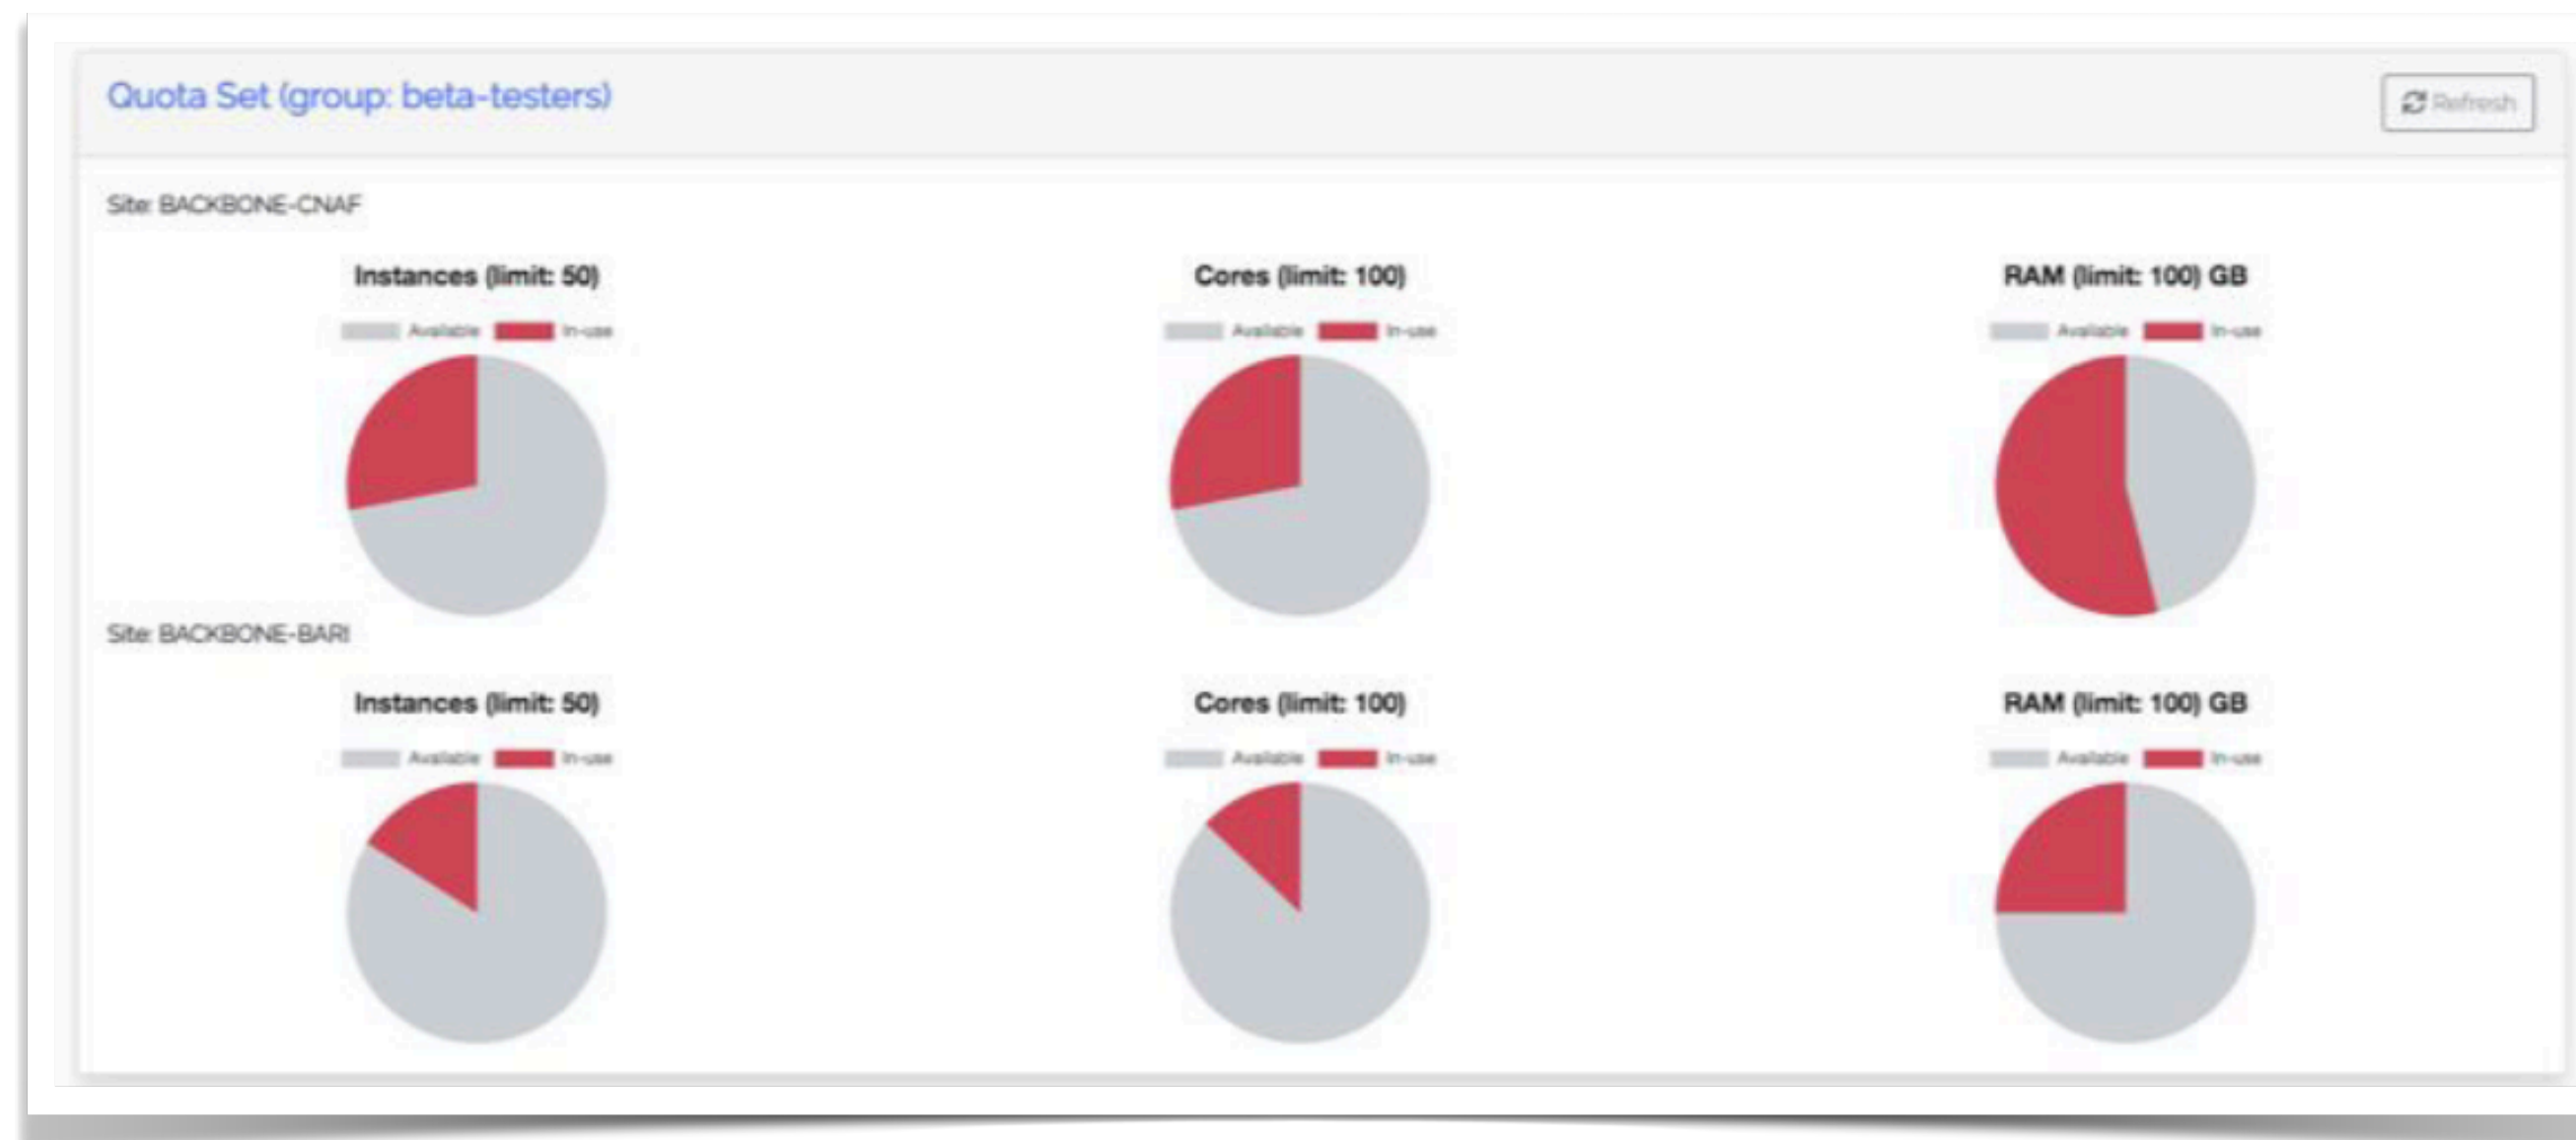
- Better Containers support.
- Add new applications.

INDIGO PaaS

SLAT (Service Level Agreement Tool) is replacing the old SLAM (from Cyfronet).

Improve resources federation.

New features are/will be available, e.g. per-group SLAs, user limits and quota management.



CREATE SLA

General information:

Service: CLOUD-CNAF_6B5EE096-D9FC-4926-A965-25F49098AEFB

Customer Group: ML-INFN

Effective from: 10/03/2021

Expiration Date: 31/12/2022

Allocated resources:

Number of Virtual Machines (if applicable): 50

Total number of vCPUs (if applicable): 256

Total amount of RAM in GB (if applicable): 921

Total number of Public IPs (if applicable): 10

Total amount of storage in GB (if applicable): 1000

User Limits (optional):

Resource type	Limit
Instances	2

SEND

Galaxy deployment

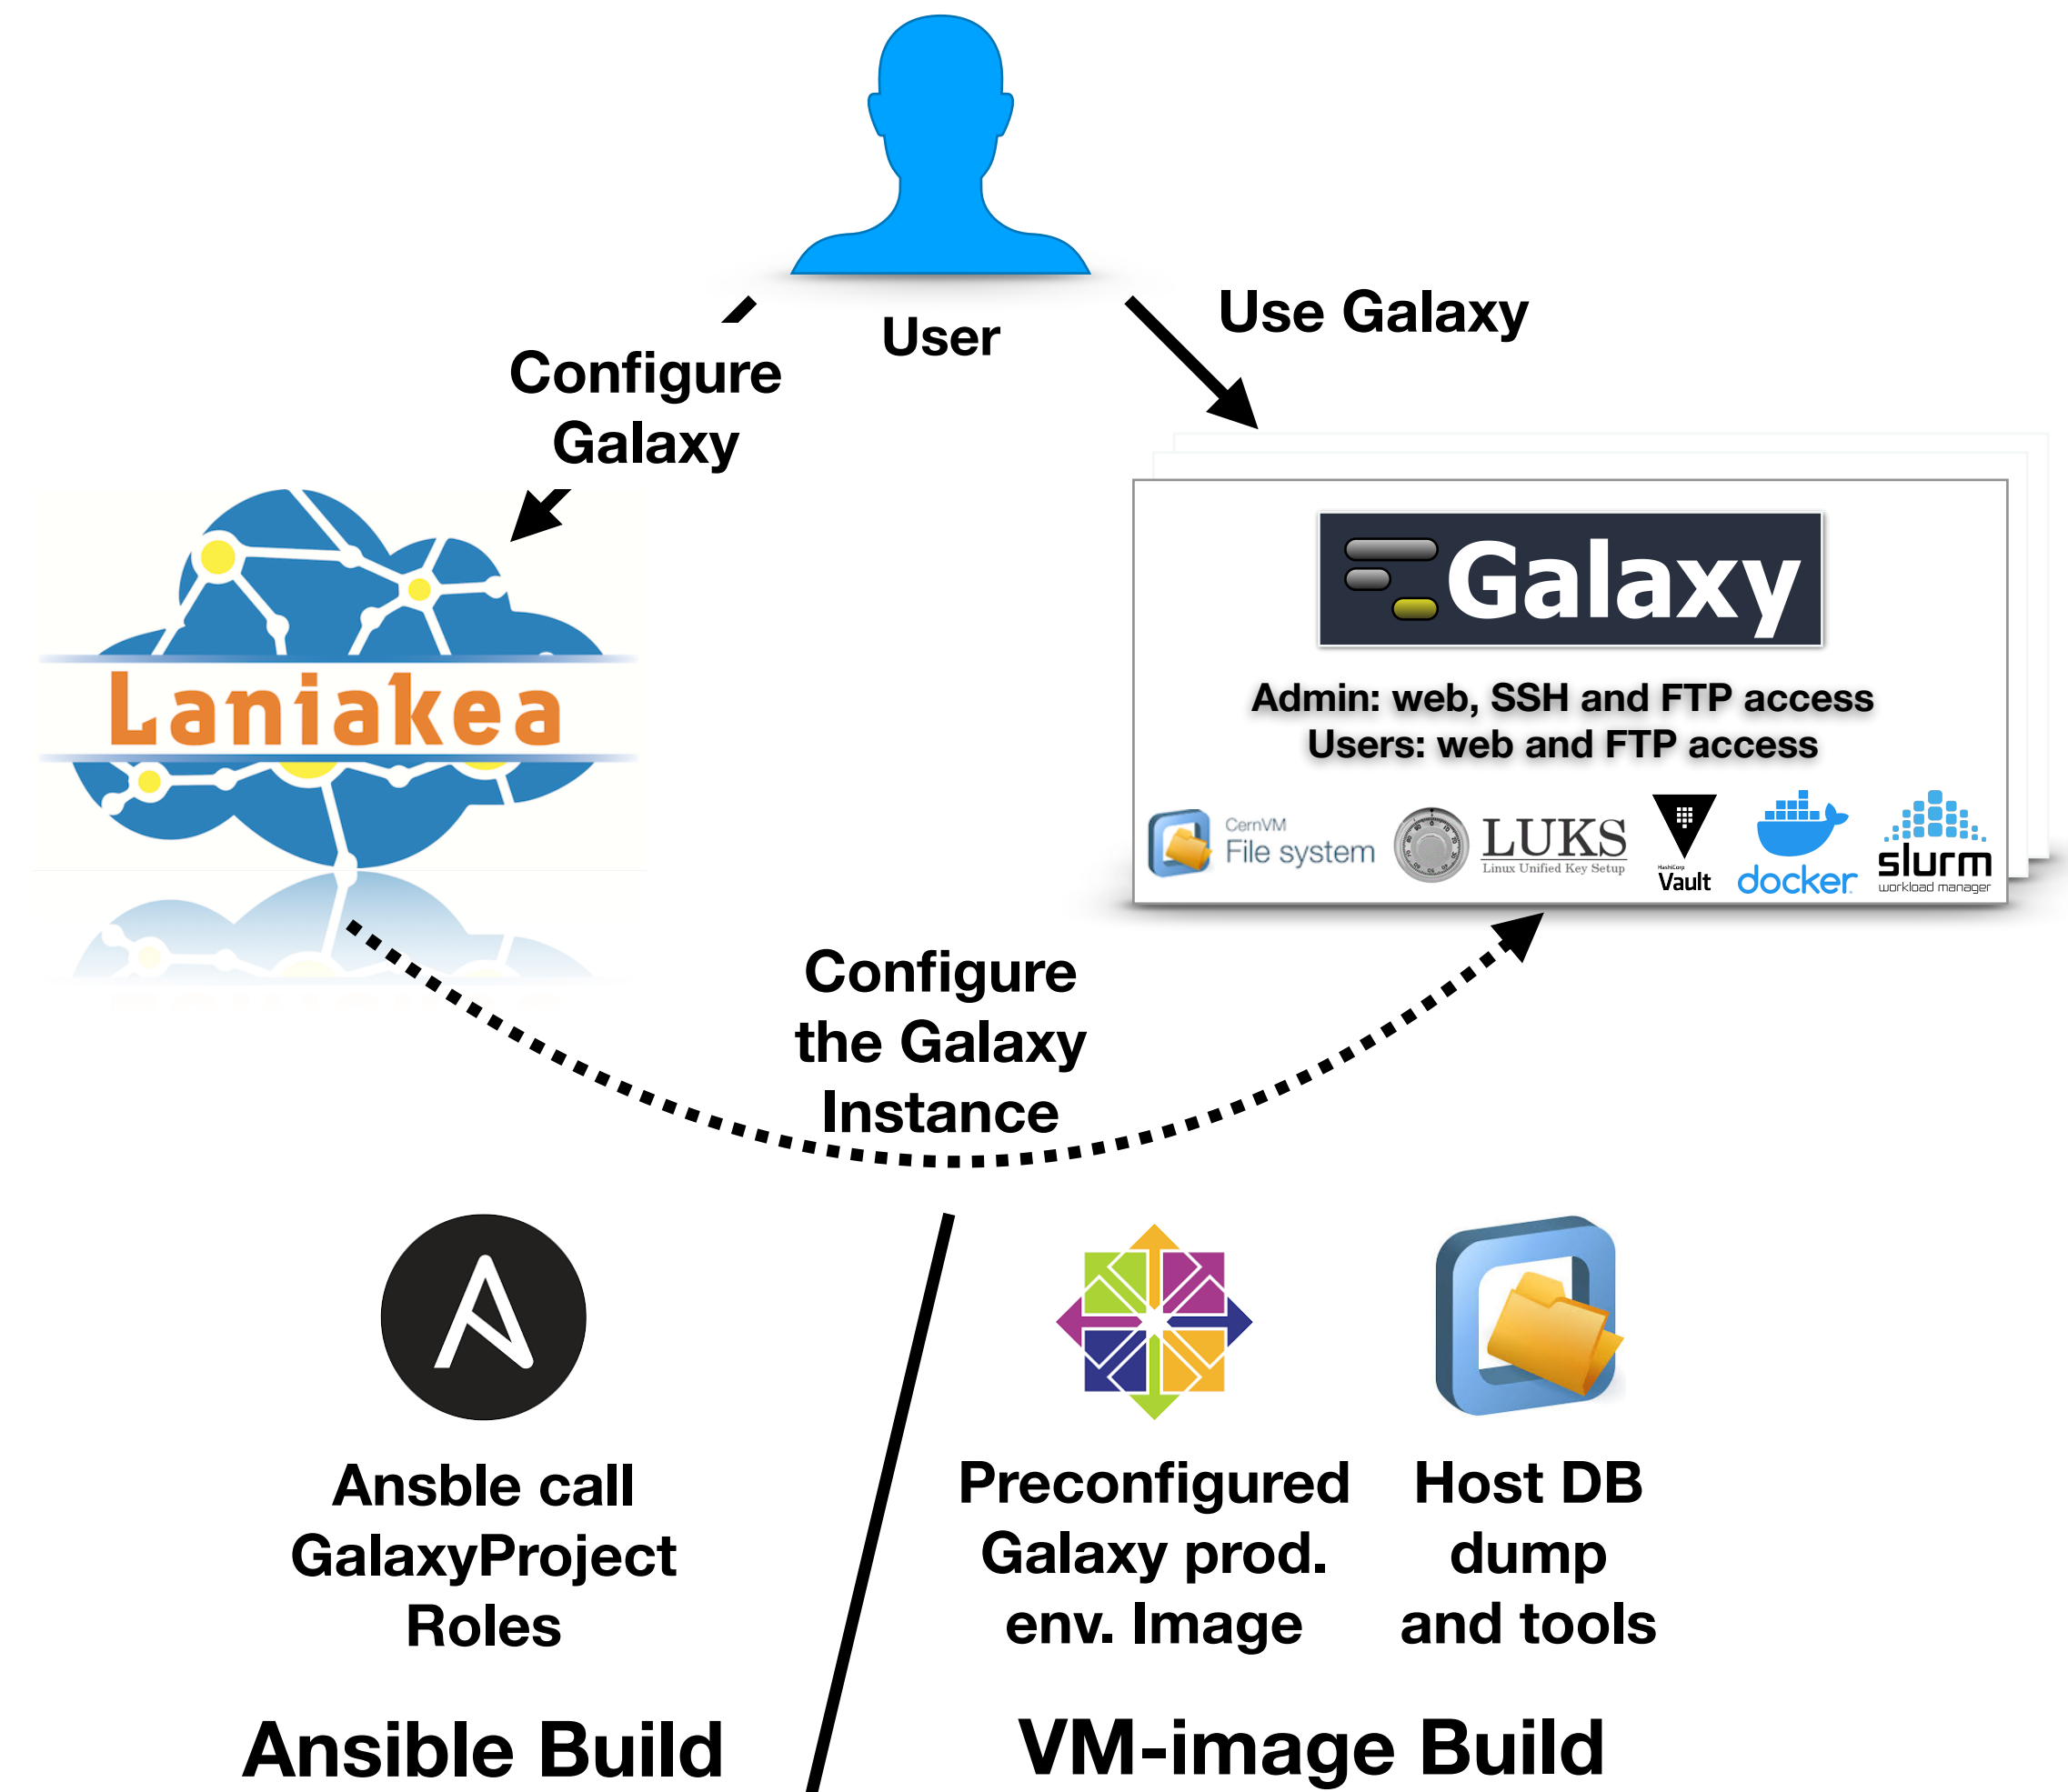
Galaxy flavours deployment through Ansible and VM-image has been completely reworked.

GalaxyProject ansible roles are used to deploy Galaxy.

CVMFS volume as a tools repository for the deployment of the Galaxy instances from VM-images.

New VM-images will be automatically created using Jenkins and Hashicorp Packer.

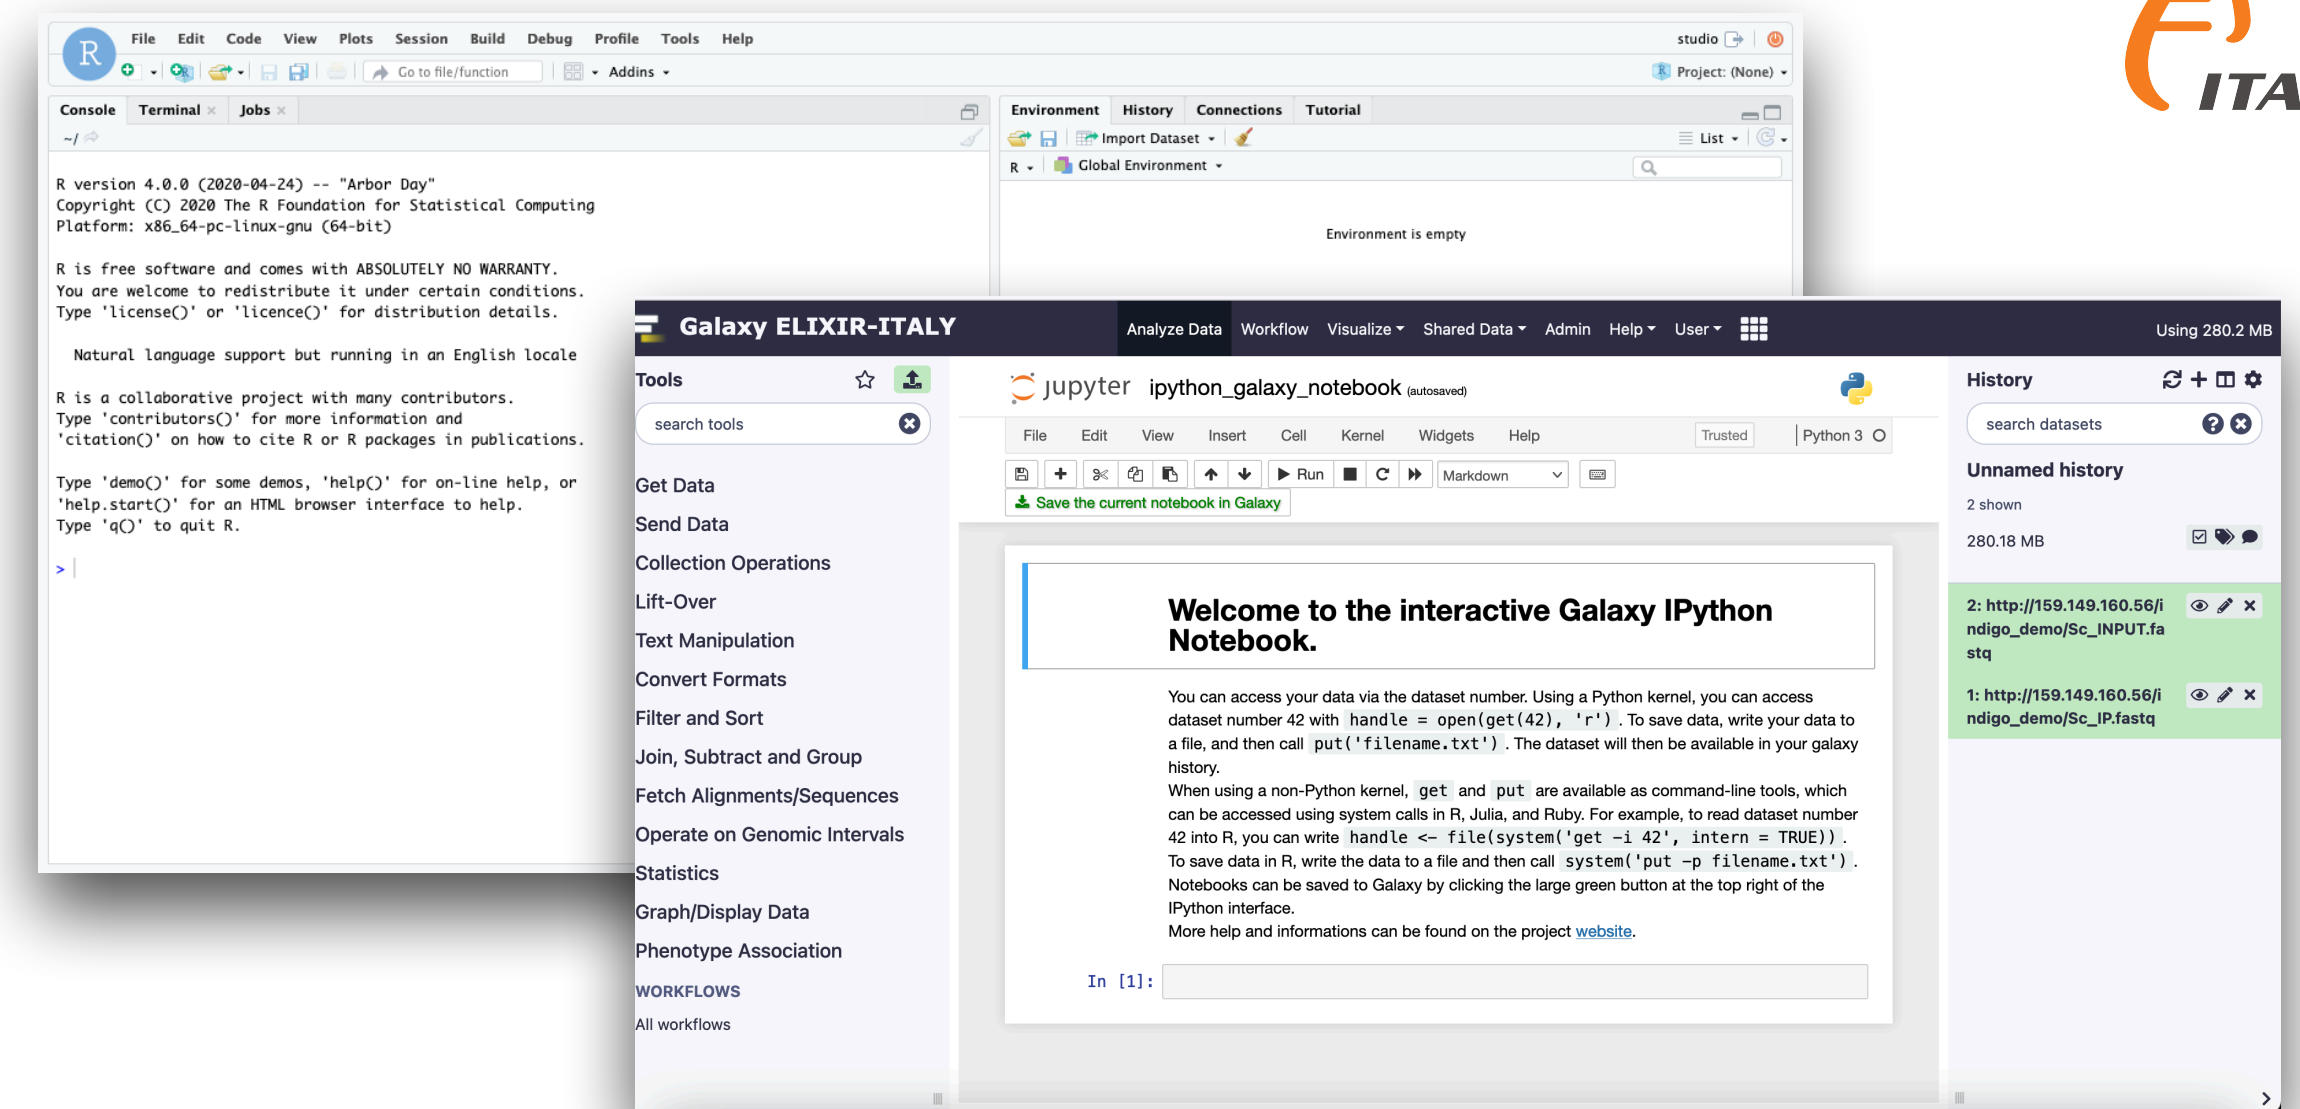
Every python2 script we were using have been ported to python3.



Users feedback

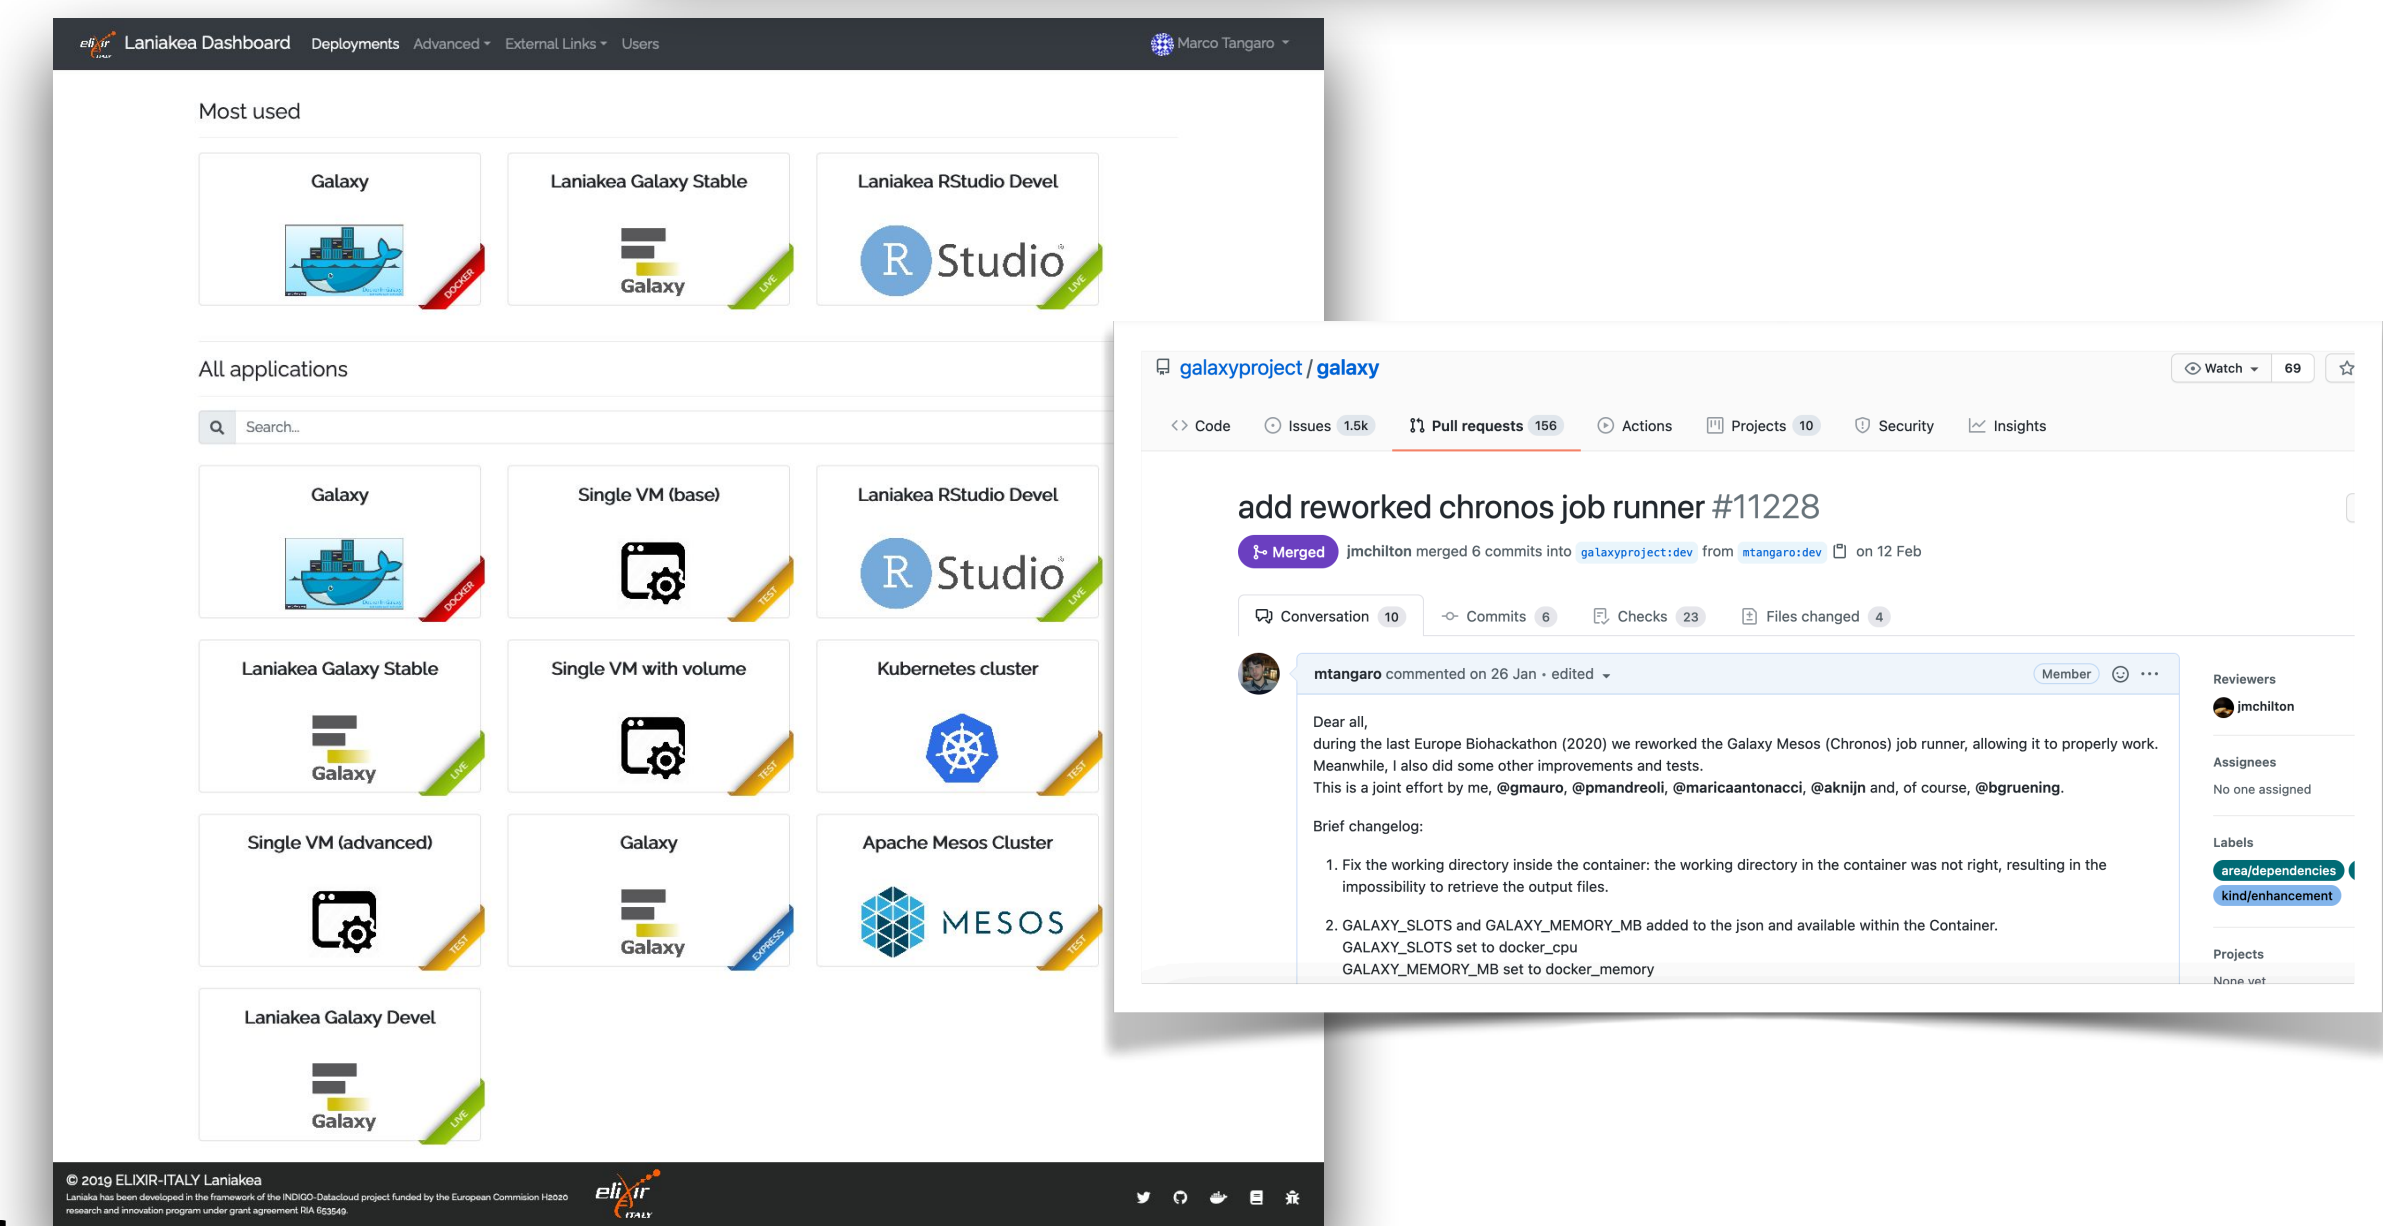
Add RStudio and Jupyter as standalone applications and integrated in Galaxy (backport to current production version on-going).

-> Requested for training and development purposes.



Add support for Galaxy with Mesos or Kubernetes for tools container orchestration.

-> RCASC user case, scRNA-Seq workflow developed using only Dockerized tools.



Thanks for your attention

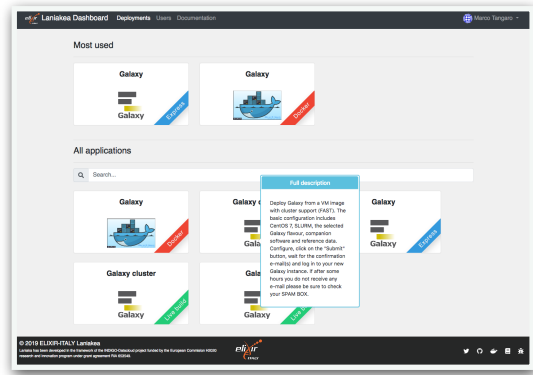
Please visit our website: <https://laniakea-elixir-it.github.io/>

Or contact us to laniakea.helpdesk@gmail.com



BACKUP

Laniakea




User friendly interface to configure both virtual hardware and software environments.



Production ready Galaxy instances and several Galaxy flavours pre-loaded with useful tools. Full admin rights over deployed instances.




Reference data volume shared across multiple Galaxy instances through read only CernVM-FS.



Support for Galaxy deployment from: VM images, Ansible roles, Docker containers.



Support for single node, cluster and elastic cluster virtual hardware.



Storage encryption to protect data privacy.

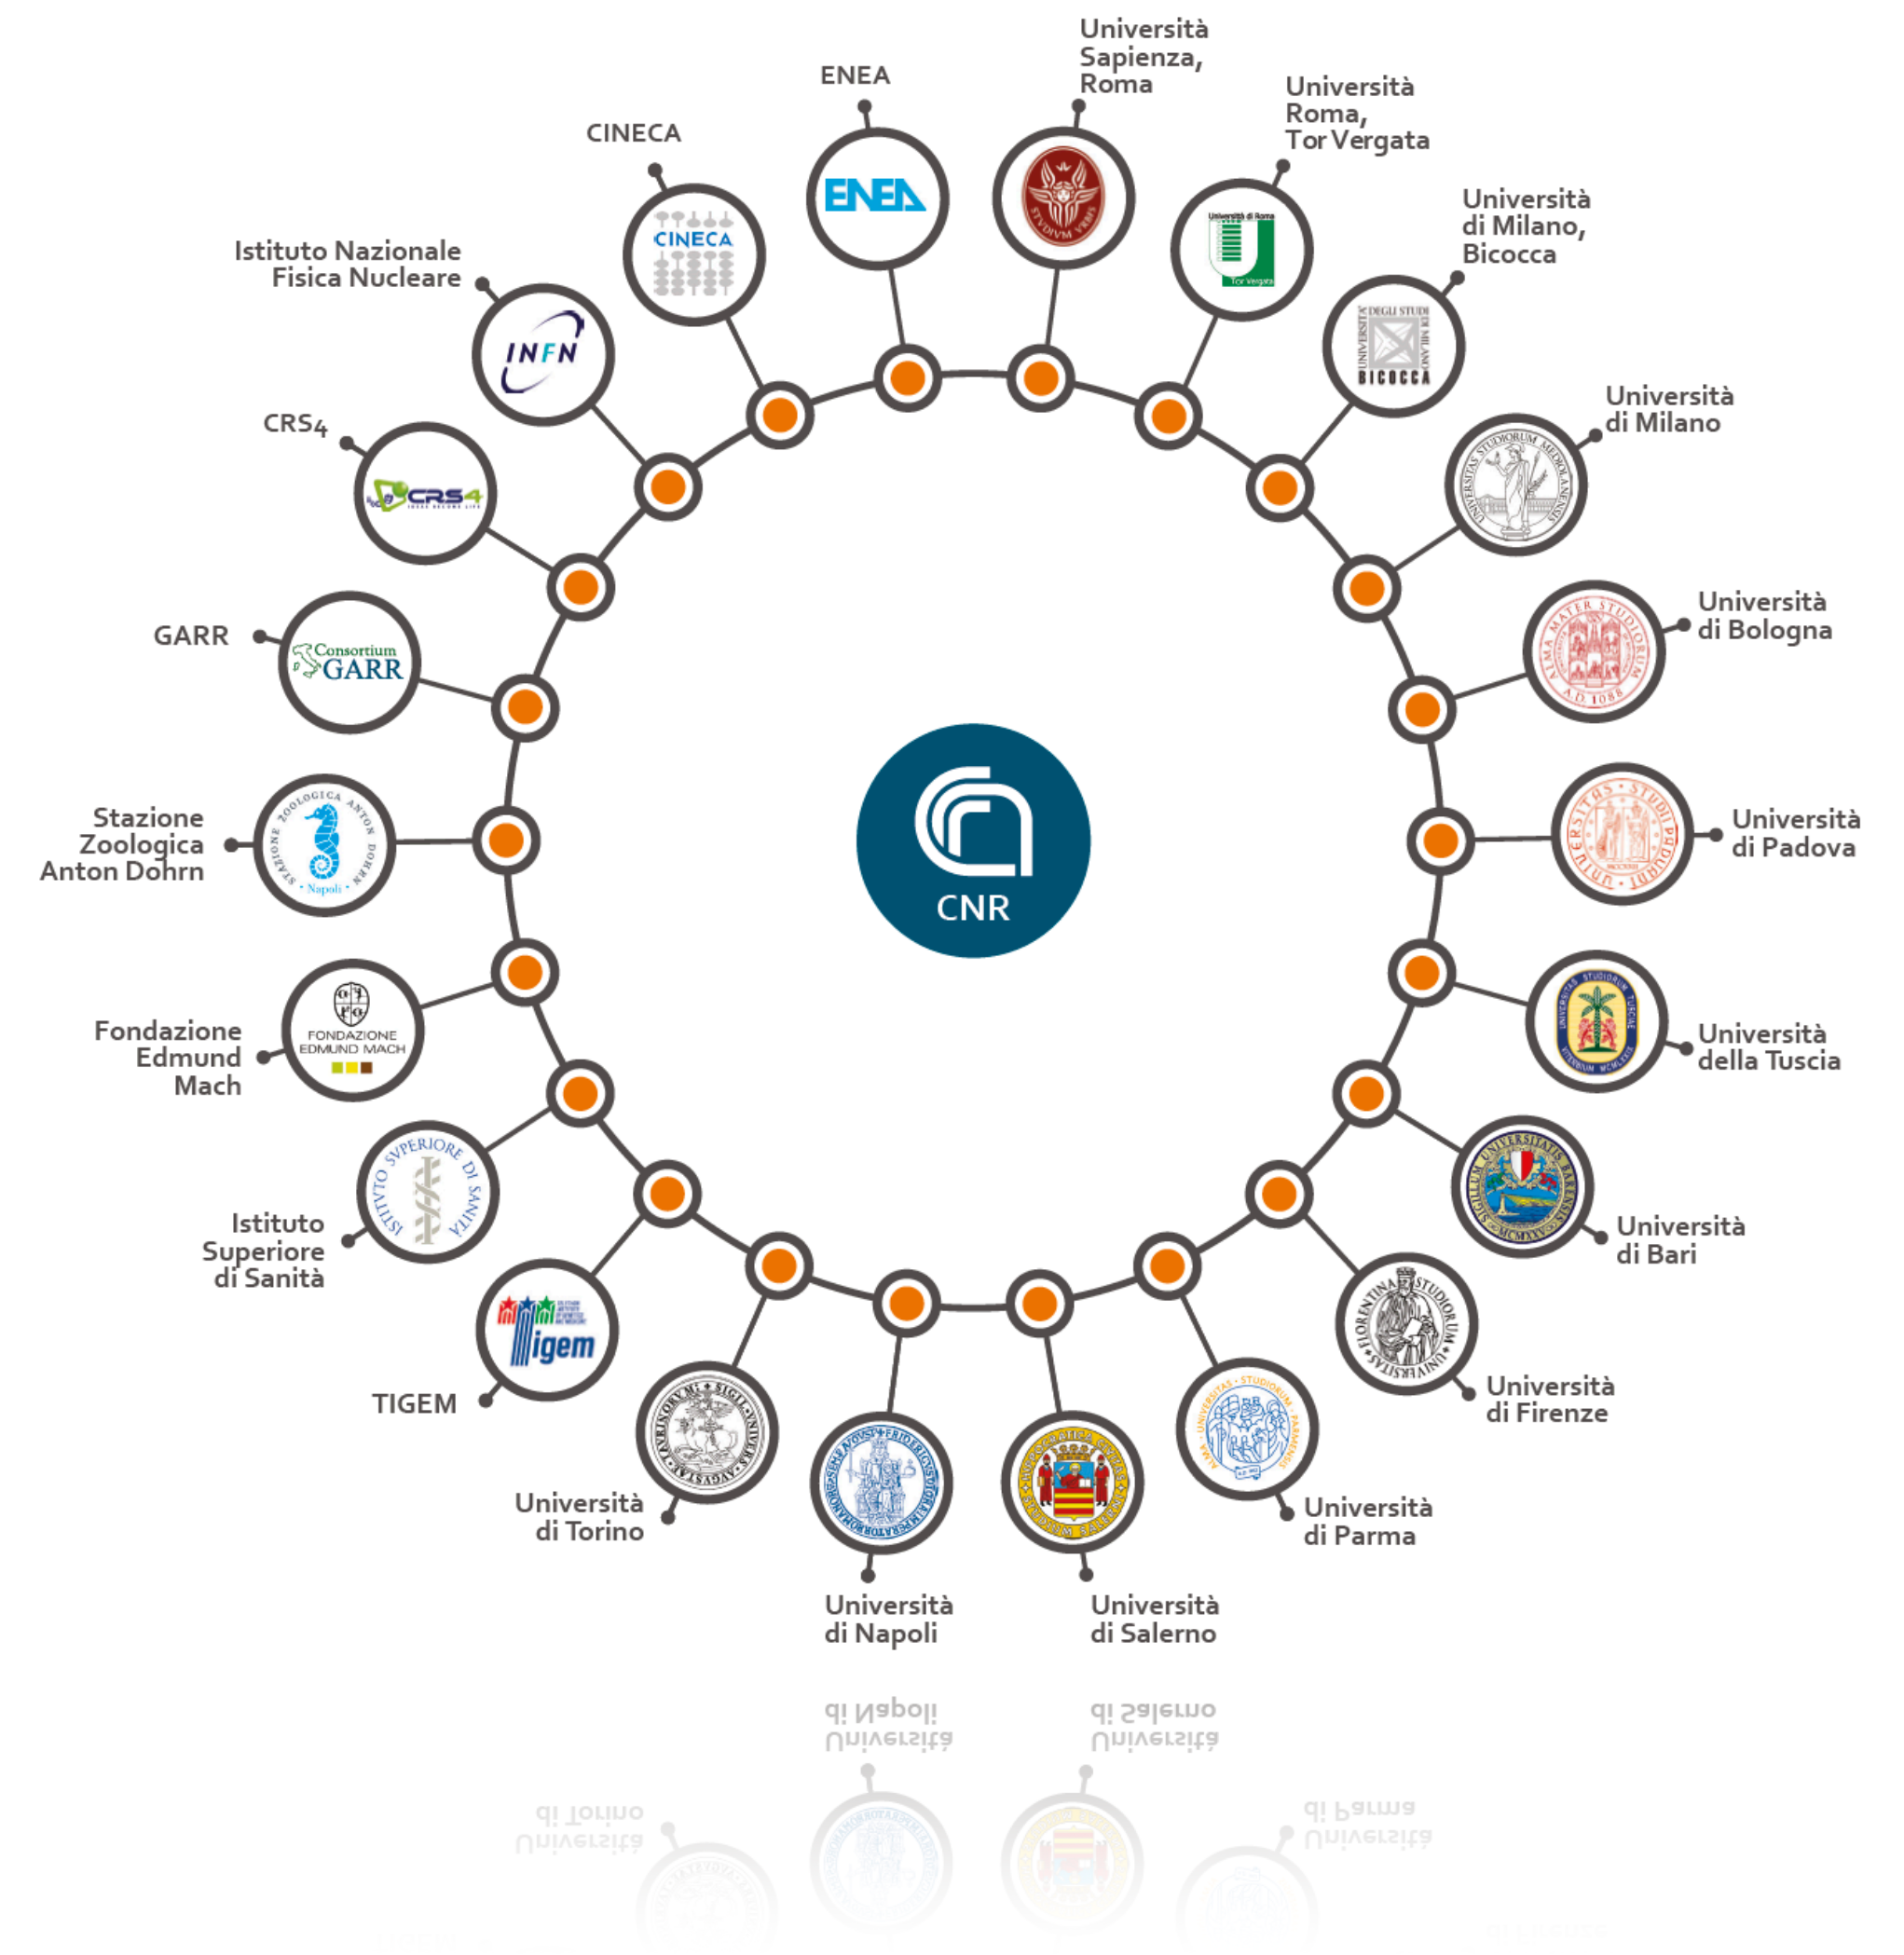
Laniakea@ReCaS service



ELIXIR-Italy partners are actively involved in the service development and/or also contribute with cloud resources.

A Laniakea service is in production for ELIXIR-ITALY partner but also for ELIXIR and external users.

The ELIXIR-ITALY Laniakea@ReCaS Call offers access to Cloud resources to be used for the deployment of on-demand Galaxy instances.

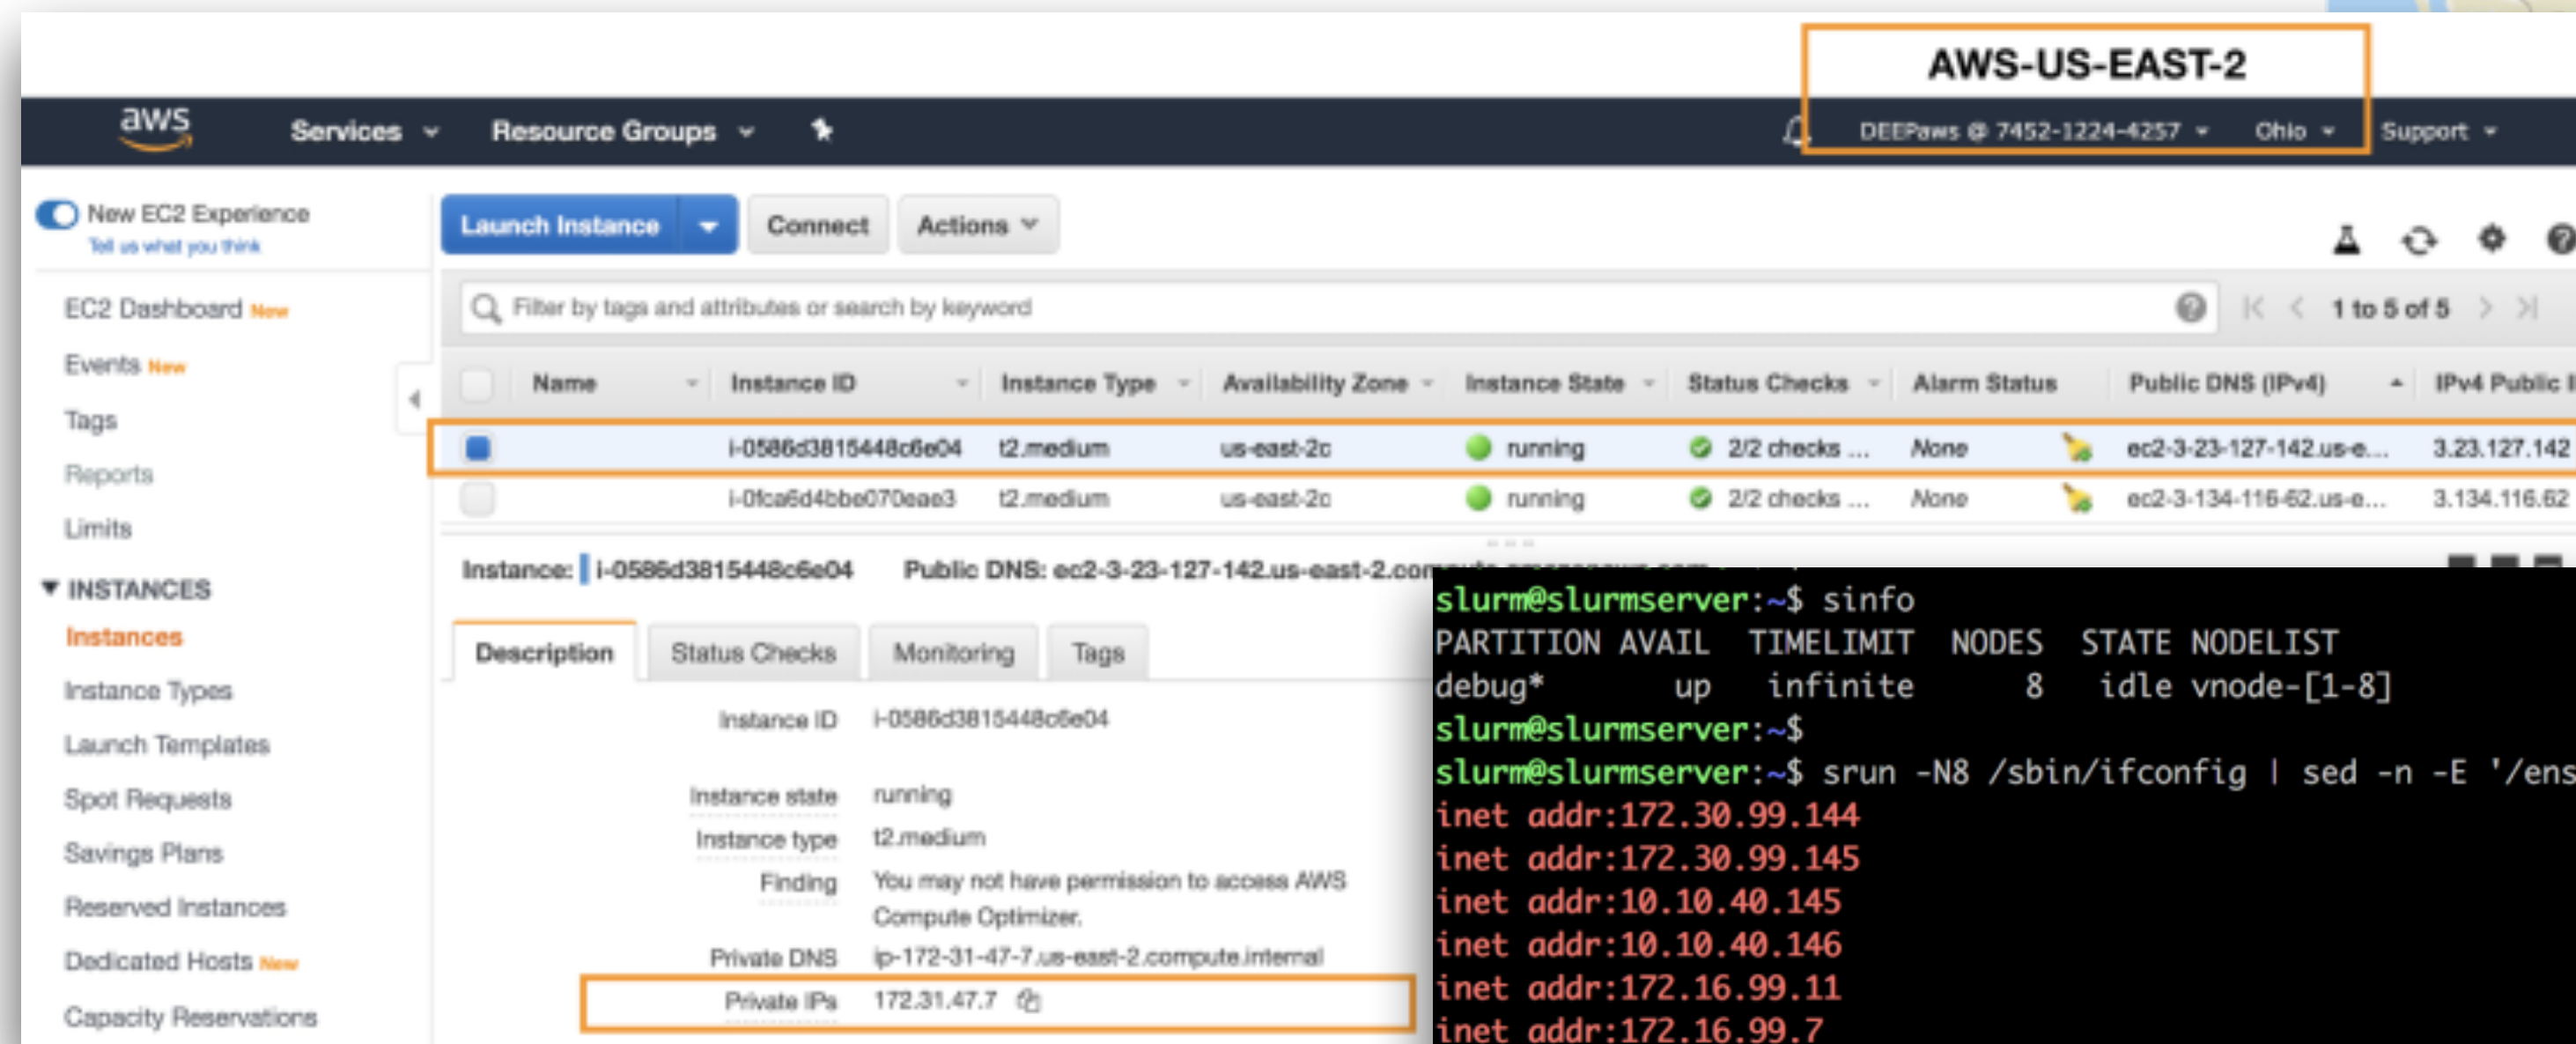


PaaS improvements

Scaling accross multiple cloud providers:

SLURM cluster with nodes on 4 sites including AWS

- INFN-BARI
- INFN-CNAF
- IFCA-LCG2
- AWS-US-EAST-2 (Ohio)



Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status	Public DNS (IPv4)	IPv4 Public IP
slurmserver	i-0586d3815448c6e04	t2.medium	us-east-2c	running	2/2 checks ...	None	ec2-3-23-127-142.us-east-2.compute.amazonaws.com	3.23.127.142
	i-0fca6d4bbe070eae3	t2.medium	us-east-2c	running	2/2 checks ...	None	ec2-3-134-116-62.us-east-2.compute.amazonaws.com	3.134.116.62

Instance: i-0586d3815448c6e04 Public DNS: ec2-3-23-127-142.us-east-2.compute.amazonaws.com

Description	Status Checks	Monitoring	Tags
Instance ID	i-0586d3815448c6e04		
Instance state	running		
Instance type	t2.medium		
Finding	You may not have permission to access AWS Compute Optimizer.		
Private DNS	ip-172-31-47-7.us-east-2.compute.internal		
Private IPs	172.31.47.7		

```
slurm@slurmserver:~$ sinfo
PARTITION AVAIL TIMELIMIT NODES STATE NODELIST
debug* up infinite 8 idle vnode-[1-8]
slurm@slurmserver:~$
slurm@slurmserver:~$ srun -N8 /sbin/ifconfig | sed -n -E '/ens[3-4]leth[0-1]/[n;p;}' | grep -Eo 'inet addr:[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}'
inet addr:172.30.99.144
inet addr:172.30.99.145
inet addr:10.10.40.145
inet addr:10.10.40.146
inet addr:172.16.99.11
inet addr:172.16.99.7
inet addr:172.31.47.7
inet addr:172.31.40.223
slurm@slurmserver:~$
```


Deployment improvement

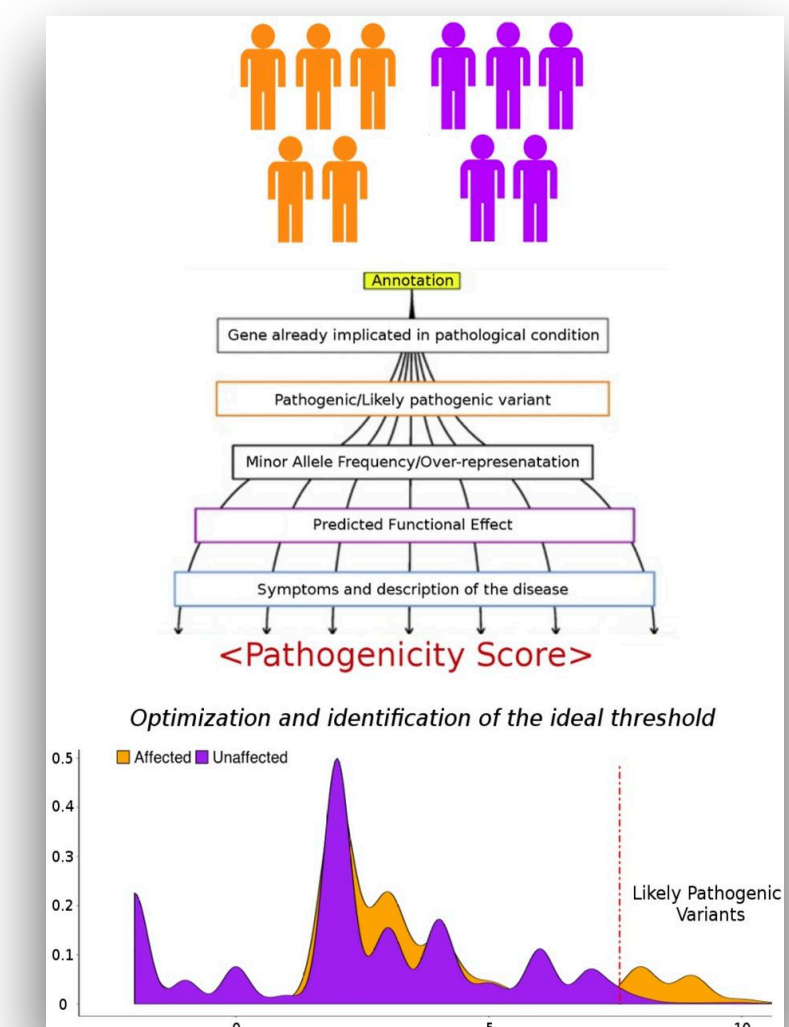
Galaxy flavours deployment through VM-image snapshots has also been completely reworked, using a CVMFS volume as a tools repository for the deployment of the Galaxy instances.



Faster alignment of Laniakea flavours to the latest Galaxy releases, the Laniakea core has been modified to exploit the official GalaxyProject ansible roles.

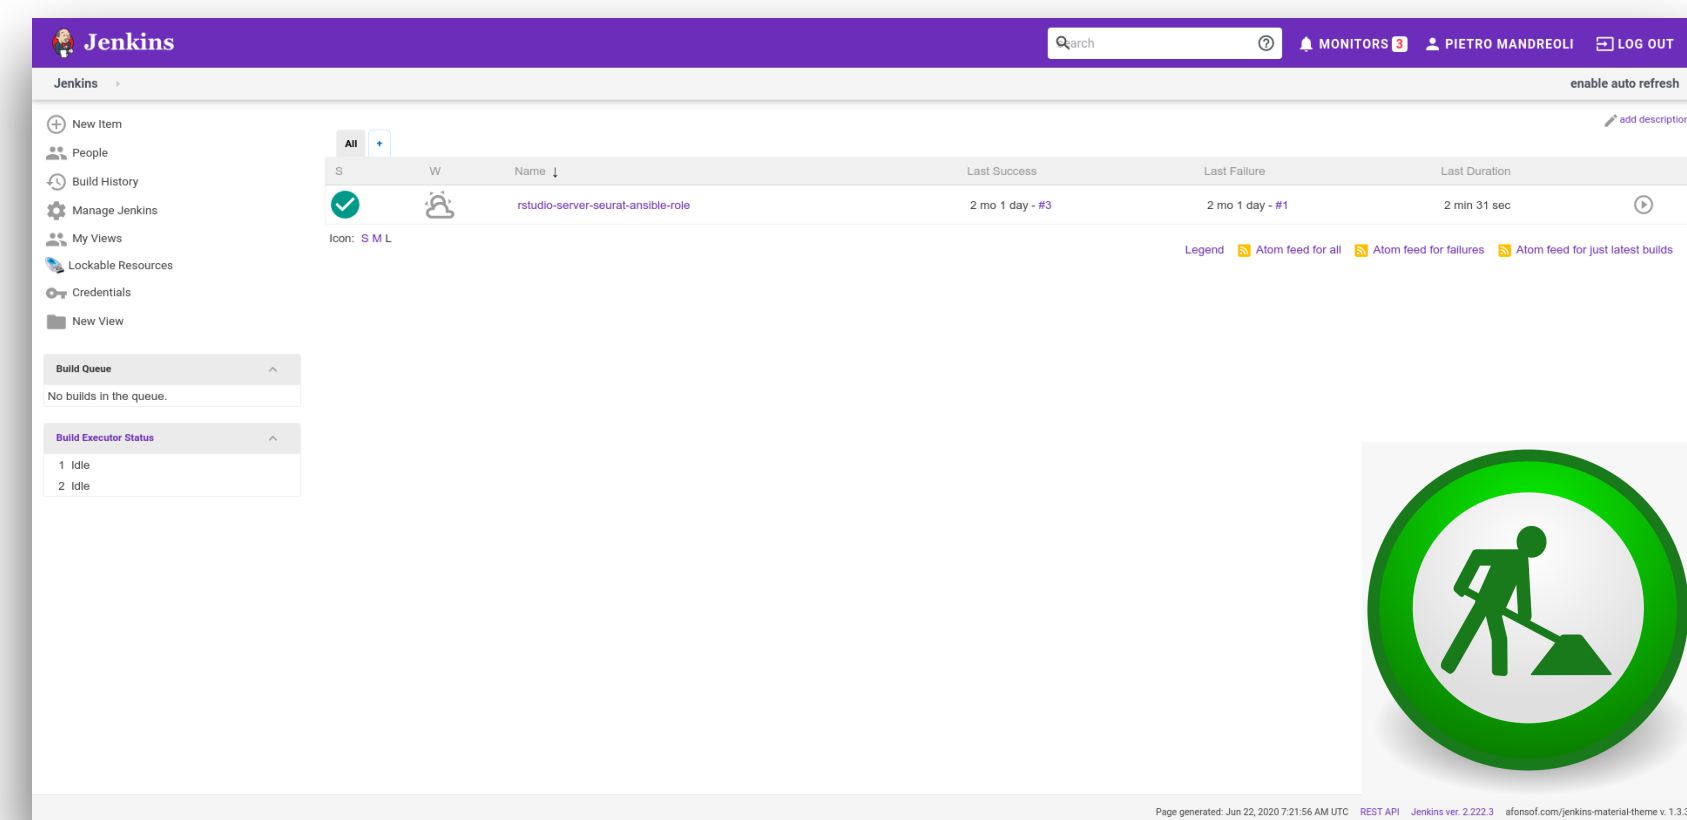
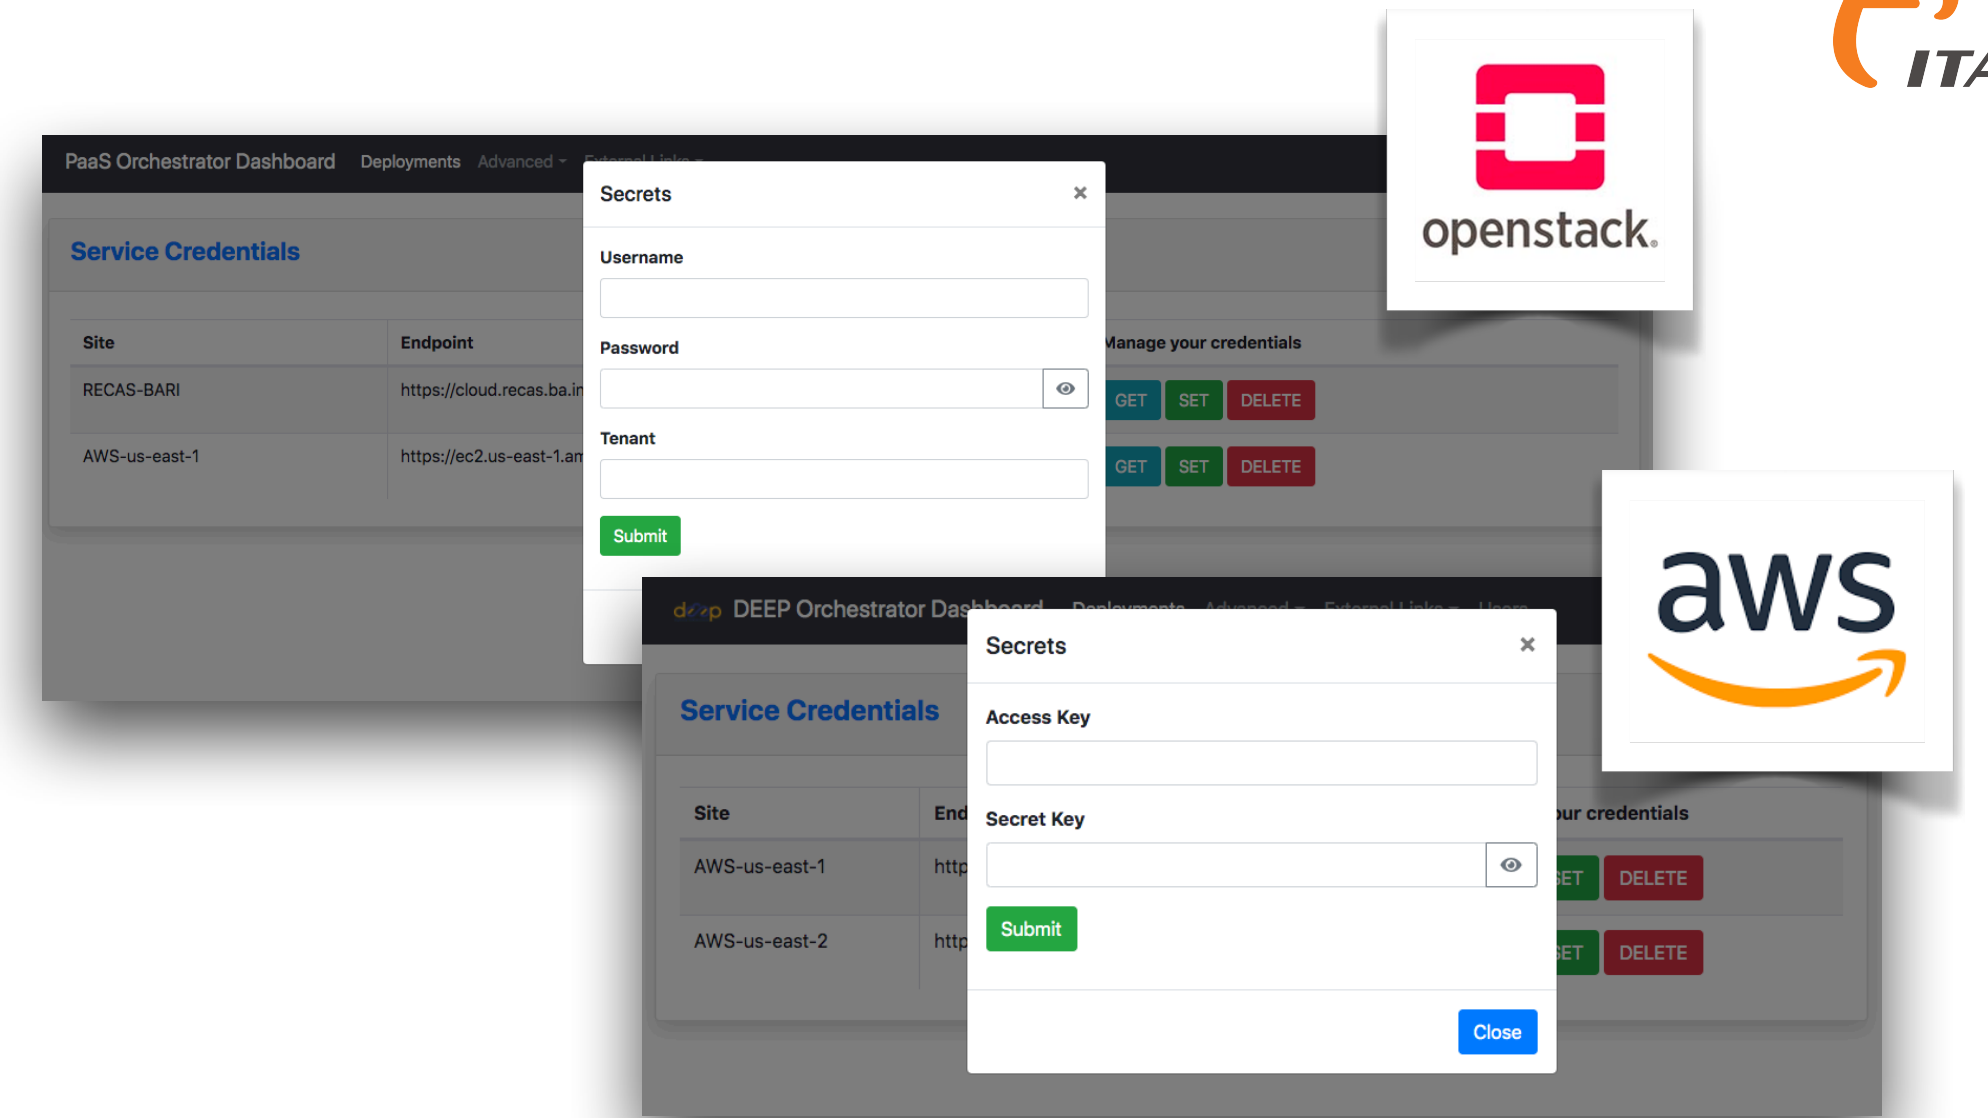
More applications beyond Galaxy incoming, like RStudio and Jupyter Notebooks.

More Flavours: VINYL highly accurate and fully automated system for the functional annotation and prioritization of genetic variants in large scale clinical studies.



Platform improvements

Improved support for clouds providing a Laniakea service, e.g., by removing the need of mandatory adoption of the AAI layer integration by cloud providers.



Software Quality Assurance: Continuous Integration framework, based on Jenkins for automatic flavour testing, update, and implementation.

Molecule will be used for Ansible roles automatic tests.