Ref. Ares(2025)5191957 - 30/06/2025

D2.5

FAIR data, reproducibility, and provenance

Tommaso Gorni, Cristiano Malica, Nicola Marzari, Fabio Affinito

Due date of deliverable	30/06/2025 (**month 30**)
Actual submission date	30/06/2025
Final version	30/06/2025
Lead beneficiary	CINECA (participant number 7)
Dissemination level	PU - Public

Document information

Project acronym	MaX
Project full title	Materials Design at the Exascale
Research Action Project type	Centres of Excellence for HPC applications
EuroHPC Grant agreement no.	101093374
Project starting/end date	01/01/2023 (month 1) / 31/12/2026 (month 48)
Website	http://www.max-centre.eu
Deliverable no.	D2.5
Authors	Tommaso Gorni, Cristiano Malica, Nicola Marzari, Fabio Affinito
To be cited as	T. Gorni et al. (2025): FAIR data, reproducibility, and provenance. Deliverable D2.5 of the HORIZON-EUROHPC-JU-2021-COE-01 project MAX (final version as of 29/06/2025). EC grant agreement no: 101093374, CINECA, Consorzio Interuniversitario Cineca.

Disclaimer

Contents

Executive Summary

The present document is a report on the deployment of the MAX model for FAIR data, reproducibility, and provenance, as detailed in task T2.6 of work package 2 (WP2), titled "Exascale workflows and extreme data for materials". FAIR data-management principles are enforced in MAX by using Materials Cloud as its main data-sharing and dissemination platform (Sec. 1). This is one of the two repositories for materials science data approved by the European Commission via Open Research Europe.

The Materials Cloud infrastructure is greatly enhanced in terms of data sovereignty and reliability via the deployment of local, EU-based backups (Sec. 2.1), a failover service for the frontend of its Archive section (Sec. 2.2) and an upgrade of its backend (Sec. 2.3). Further improvements of the present infrastructure and additional services to be deployed in the future are discussed in Sec. 4.

1 The MaX FAIR data model and the Materials Cloud Archive

The MAX data management model relies on Materials Cloud [1], an open platform *built to enable the seamless sharing and dissemination of resources in computational materials science, offering educational, research, and archiving tools; simulation software and services; and curated and raw data* [2]. The Materials Cloud Archive section is an open-access, moderated repository for research data in computational materials science, which implements and enforces the FAIR principles [3] for data management, in particular:

1. It allows researchers worldwide to upload, publish and download data free of charge;
2. It provides globally unique and persistent digital object identifiers (DOIs) for every record;
3. Metadata can be harvested in a number of machine-readable formats;
4. Data are accessible through an interactive front-end interface as well as a REST API.

Moreover, the Materials Cloud Archive provides its own guidelines and templates for good data management practices [4], and it is indexed by Re3data [5] and EUdat B2find [6]. A snapshot of how the data deposited on the Materials Cloud Archive can be browsed, manipualted and visualized is presented in Fig. 1.

Prior to the beginning of MAX phase-3, the Materials Cloud Archive data were hosted at the Swiss National Supercomputing Centre (CSCS). One of the goals of WP2, specifically of T2.6, is the creation of persistent backups of the Archive in EU HPC centres. Such an achievement can serve as a basis for two further developments: (1) the deployment of a completely independent front-end interface, which can act a Materials Cloud failover and (2) the possibility for the users of the HPC centre to access and use data locally, without the need to download them from the central repository.

This first enabling target has already been reached with the creation of nightly backups of the Materials Cloud Archive towards the CINECA's (Italy, EU) and Jülich's (Germany, EU) supercomputing centres, as detailed in Sec. 2.1. Moreover, a read-only failover of the Archive has been deployed to production at the CINECA site, as shown in Sec. 2.2. Finally, an upgrade of the Materials Cloud backend (from the Invenio Framework [7] to the InvenioRDM application [8]) will go to production in the second half of 2025 at CSCS. The new backend is also currently being tested for the CINECA failover in a development environment (Sec. 2.3) so to avoid any discontinuity in the failover service.

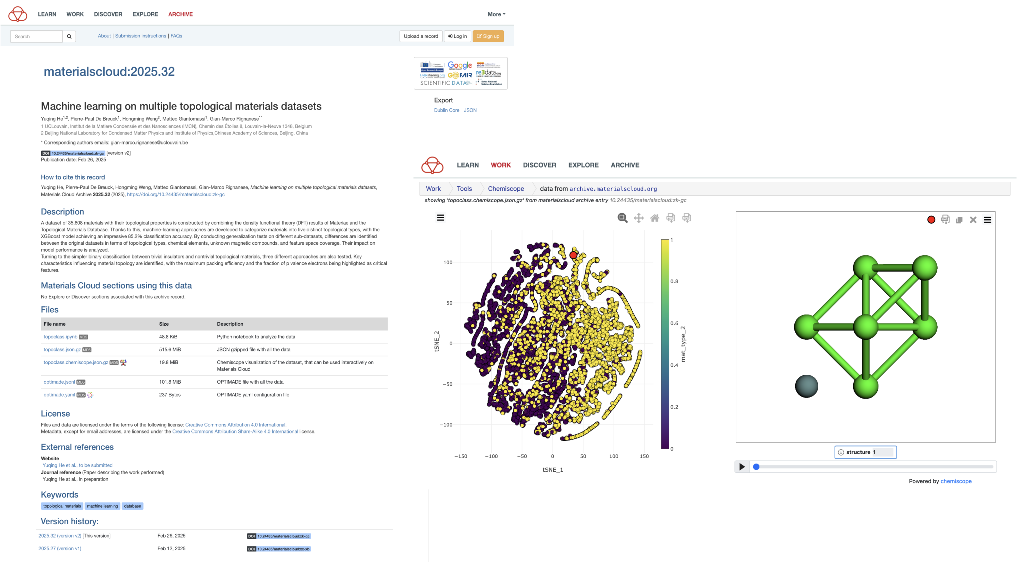

Figure 1: In-the-cloud manipulation and visualization of FAIR data deposited in the Materials Cloud Archive.

2 Upgrade of the MaX FAIR data resources

In March 2023, cloud resources have been allocated on CINECA's Ada Cloud [9] to support the upgrade of the Materials Cloud infrastructure. The allocated tenant, represented in Figure 2, has been configured to host two virtual machines (VM) of 4 virtual CPUs and 30GiB of RAM each, one storage volume of 10 TiB, and two floating IPs (FIPs). The storage volume hosts the Materials Cloud Archive data backup, while the (production-) VM hosts the interactive front-end interface of the Materials Cloud Archive. The interactive interface is exposed to the external network as a failover website through a public FIP, associated to the production VM. A second VM (development-VM) and FIP have been allocated as a development environment, which is used to perform backend upgrades without affecting the production. The details of these deployments follow in the next sections.

2.1 Nightly backups to EU HPC centres

Nightly backups occur between the primary Materials Cloud Archive, hosted at CSCS in Lugano (Switzerland), and the Cinder volume at CINECA. The backup is handled by a cronjob residing on the CSCS server, which copies every night the latest uploads from the CSCS object storage to CINECA's Cinder volume. The CINECA's volume is formatted as an *ext4* file system and mounted on the production-VM. The copy process is handled by *rclone* [10], which can deal with copies from object to file-system storage. The relevant portion of the Bash script called by the cronjob is reported in Listing 1.

```
#!/bin/bash
[...]

DESTINATION=${DEST_REMOTE}:${DEST_BUCKET}/${DEST_BUCKET_FOLDER}
```

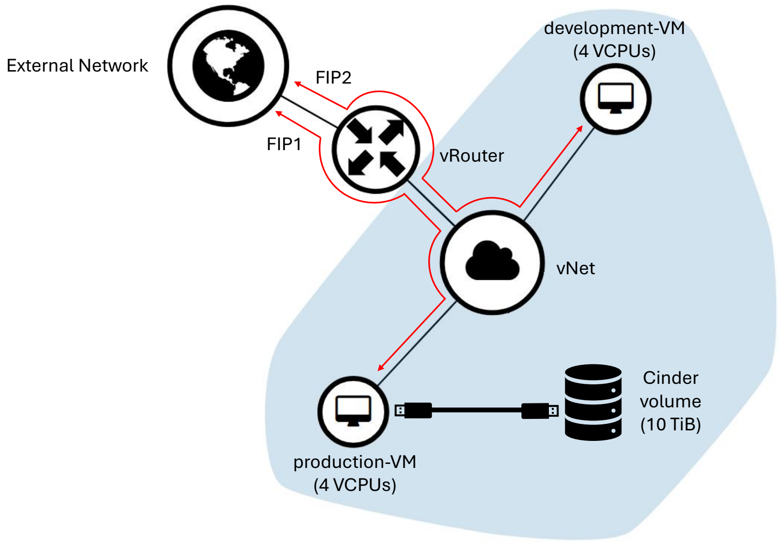

Figure 2: Sketch of the configuration of the OpenStack tenant on ADA Cloud used to host the Materials Cloud Archive failover at CINECA. The two VMs are part of the same subnet (vNet), while a virtual router handles the communications to and from the external network.

```
DESTINATION_PREVIOUS=${DEST_REMOTE}:${DEST_BUCKET}/${DEST_BUCKET_FOLDER_PREVIOUS
    }

# Size of DESTINATION before backup
# --------------------------------
echo
echo "OVERALL SIZE OF DATA AT ${DESTINATION} BEFORE BACKUP:"
echo $(rclone ls ${DESTINATION} | awk 'BEGIN{a=0}{a+=$1}END{print a
    /1024./1024./1024., "GB"}')
echo

# Copy files to DESTINATION
# -------------------------------
echo "COPYING from CSCS container ${SOURCE_BUCKET} of REMOTE ${SOURCE_REMOTE} to
     ${DEST_BUCKET} of REMOTE ${DEST_REMOTE}..."
rclone copy -v ${SOURCE_REMOTE}:${SOURCE_BUCKET} ${DESTINATION} --backup-dir=${
    DESTINATION_PREVIOUS} --suffix=".${TIMESTAMP}"
rc_dest=$?

echo "RE-CHECKING..."
rclone check -v --one-way ${SOURCE_REMOTE}:${SOURCE_BUCKET} ${DESTINATION}
rc_dest_check=$?
echo "Done!"
echo

# Check for changes
# -------------------------------
echo "CHECKING WHETHER ANY FILES CHANGED CONTENT"
prev_dest=$(rclone ls ${DESTINATION_PREVIOUS} | tee -a /dev/stderr | grep ${
    TIMESTAMP} | wc -l)
echo "${prev_dest} changed files at ${DEST_BUCKET}/${DEST_BUCKET_FOLDER_PREVIOUS
    }"

# Size of DESTINATION after backup
# --------------------------------
```

```
echo
echo "OVERALL SIZE OF DATA AT ${DESTINATION} AFTER BACKUP:"
echo $(rclone ls ${DESTINATION} | awk 'BEGIN{a=0}{a+=$1}END{print a
    /1024./1024./1024., "GB"}')
echo
```

Listing 1: Excerpt from the *rclone*-based Bash script used to perform a copy of the Materials Cloud Archive entries from CSCS object storage to CINECA's file-system storage. The script is called by a cronjob running on a CSCS server and is executed every night.

From August 2023, when the CINECA's failover entered the production phase, to May 2025, the overall disk-space occupation has increased from 1.2 TiB to 2.2 TiB. It must be noted that this is an overestimate of the total Archive size, since a small fraction of redundant data are copied every night by the *cronjob*. The copy process can be hence optimized in the future to avoid filling up the 10-TiB Cinder volume prematurely. Starting from 2025, a similar backup is performed towards a S3 storage hosted at the Jülich Supercomputing Center cloud services [11]. The *rclone* script has been adapted to perform a copy from S3 storage to S3 storage, the changes with respect to listing 1 are reported in listing 2.

```
#!/bin/bash
[...]
DESTINATION=${DEST_REMOTE}:${DEST_BUCKET_FOLDER}
DESTINATION_PREVIOUS=${DEST_REMOTE}:${DEST_BUCKET_FOLDER_PREVIOUS}
[...]
rclone copy -v --s3-max-upload-parts=1000 ${SOURCE_REMOTE}:${SOURCE_BUCKET} ${
    DESTINATION} --backup-dir=${DESTINATION_PREVIOUS} --suffix=".${TIMESTAMP}"
[...]
```

Listing 2: Changes made to listing 1 to copy from object storage to object storage.

The deployment of these *cronjob*s guarantees two independent, daily-updated backups of the Materials Cloud Archive within the European Union, so ensuring both redundancy and sovereignty of data. As already mentioned in Sec. 1, they also serve as a starting point for a deeper upgrade of the Materials Cloud infrastructure, as shown in the next section.

2.2 The CINECA's Materials Cloud Archive failover

Currently, the deployment of the Materials Cloud website is automated by a collection of private Ansible [12] scripts (also called Ansible playbooks) targeting a given cloud architecture. These scripts, based on the OpenStack Ansible plugins, setup and configure a virtual machine as an Apache server for the Materials Cloud website. A dedicated script is devoted to the installation and configuration of the Materials Cloud Archive section, based on Invenio Framework [7]. With the help of Valeria Granata, scientist at EPFL and part of the Materials Cloud developers team, the aforementioned Ansible scripts have been adapted to run on CINECA's ADA cloud, resulting in the successful deployment in August 2023 of a Materials Cloud Archive instance based on CINECA's data backup. Public HTTPS access to the VM hosting this Materials Cloud Archive instance is granted by associating a public FIP to the VM and defining an URL via DNS association. In order to comply with CINECA's policies [13] the so-created failover website is reachable from the URL https://materialscloud-archive-failover.cineca.it. A snapshot of failover mainpage is reported in Figure 3.

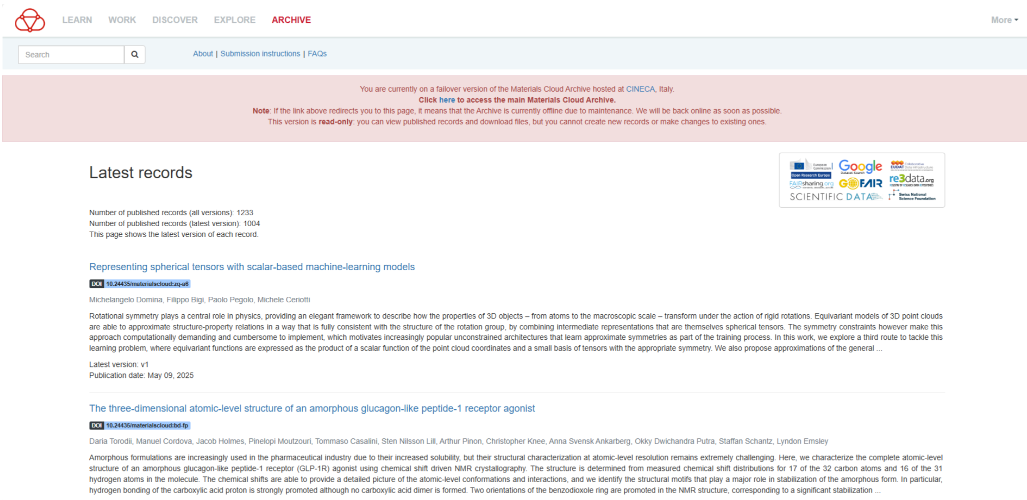

Figure 3: Homepage of the Materials Cloud Archive failover at CINECA. The red banner notifies the users that they are currently browsing a failover website.

At present, by connecting to this CINECA's failover server, users can browse and download the updated Materials Cloud Archive. The uploading of new entries has been disabled and is currently possible only on the main CSCS server, to ensure the consistency of data with the present uni-directional backups. A protocol to synchronize different Archive databases, each having uploading enabled, is under discussion, so to allow in the future the deployment of the Materials Cloud Archive as a federation of equivalent servers.

The Materials Cloud developers team is also currently implementing an automatic domain name switch in case of failure of the primary Archive, so that redirection to the CINECA's server occurs without any users' action.

2.3 Upgrade of the Materials Cloud Archive backend

The Materials Cloud project is currently migrating from the Invenio Framework [7] to the InvenioRDM backend [8]. While the former can be seen as toolbox to build one's own digital repository, InvenioRDM is a repository platform based on the Invenio Framework and Zenodo [14]. During the first quarter of 2025, a separate instance of the Materials Cloud Archive, based on the new InvenioRDM backend, has been deployed on the development virtual machine on CINECA's ADA Cloud (see Figure 2). A separate FIP and domain name have been associated to the development VM so that a preliminary version of the Materials Cloud Archive with the upgraded backend can be reached independently of the production one via the URL https://materialscloud-archive-failover-dev.cineca.it. A snapshot of failover mainpage with upgraded backend is reported in Figure 4.

Currently, the storage volume hosting the Archive data backup is attached only to the production VM, the one using the old backend, hence data are not yet downloadable from https://materialscloud-archive-failover-dev.cineca.it, which is displaying only a sample of metadata. As soon as the main, CSCS hosted, Materials Cloud website goes to production with the new backend, the storage volume will be mounted on the develop-

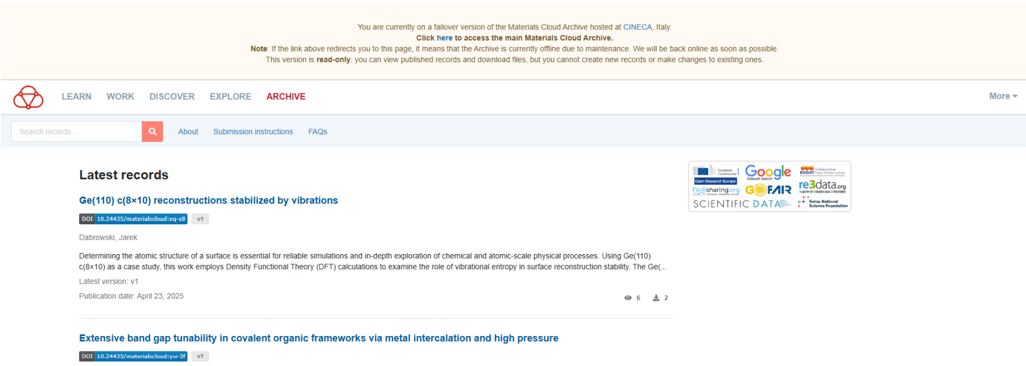

Figure 4: Homepage of the Materials Cloud Archive failover at CINECA with the upgraded InvenioRDM backend. The red banner turned yellow.

ment VM and the two domain names will be switched, so to have a seamless upgrade of the failover.

3 FAIR data use cases

Building on the AiiDA infrastructure and its integration with the Materials Cloud platform, several use cases have highlighted the potential of automated and high-throughput workflows in accelerating advanced materials discovery. Among the recent and most impactful developments within MAX has been the systematic calculation and data mining of Hubbard U and V parameters — quantities that are essential for achieving accurate electronic structure predictions, particularly in transition-metal-based compounds. This effort is especially relevant for the discovery and characterization of cathode materials for next-generation lithium-ion batteries.

The public release of the aiida-hubbard workflow marks a significant milestone in this direction. This tool enables fully automated and reproducible computation of Hubbard U and V parameters using density-functional perturbation theory. These parameters are crucial for the study of materials relevant to the energy transition, particularly rechargeable battery cathodes, where strong electronic correlations play a key role. The infrastructure developed within MAX allows for high-throughput studies of U and V across a wide range of materials. Initial benchmark efforts—detailed in Bastonero, Malica, Macke et al.[15] — have already covered over 100 Li-, Mn-, and Fe-based compounds. The workflow has since been scaled to approximately 1,000 candidate materials for Li-based cathodes, with plans to expand to tens of thousands.

A key advantage of this systematic approach is the integration with the Materials Cloud Archive, which enables structured storage and open dissemination of U and V datasets. This not only ensures long-term accessibility but also facilitates the application of data-driven methods. In particular, the calculated values are being used to train advanced machine learning models—based on equivariant architectures—that can accurately predict Hubbard parameters at negligible computational cost.

This opens the door to AI-powered interfaces [16]: we are currently developing a web-based application that will allow users to upload a crystal structure and instantly obtain predicted U and V values. This tool will provide researchers with an efficient alternative to costly first-principles calculations, accelerating the screening and optimisation of functional materials for energy storage and beyond. These developments showcase how the MAX infrastructure supports the full data lifecycle—from first-principles simulation to AI-ready datasets—offering robust, interoperable, and reusable tools for the broader materials science community.

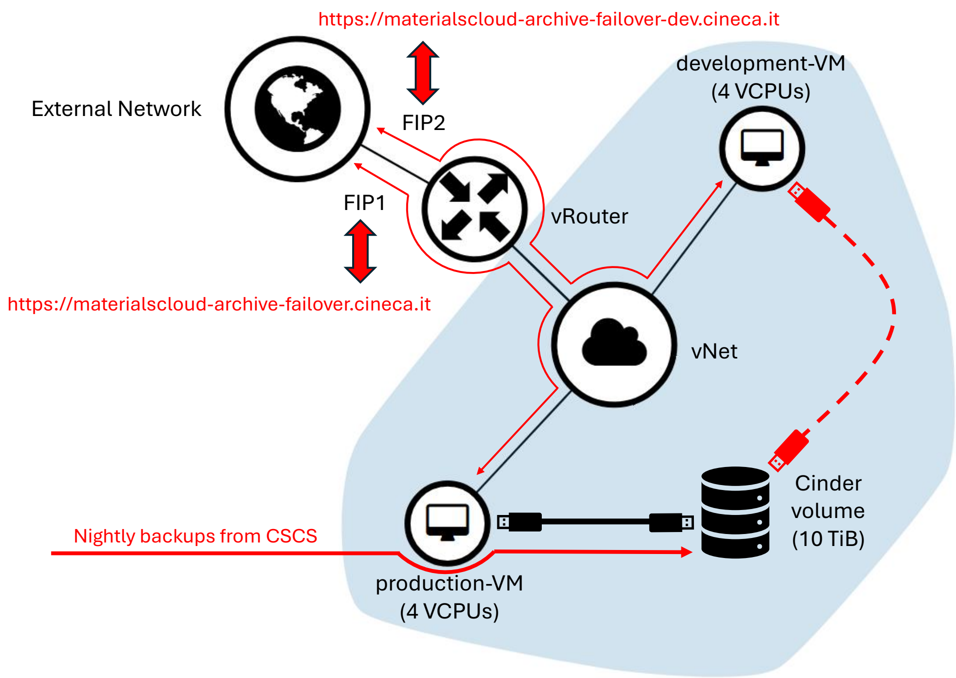

Figure 5: Graphical representation of the main deployments presented in this report and of the used resources. The dashed USB cable represents the switching to the new backend.

4 Planned improvements

To summarize, the following goals have been achieved:

- Local, nightly, backups of the Materials Cloud Archive are performed from CSCS to CINECA (Italy, EU) and Jülich (Germany, EU).
- The CINECA database is used to deploy an independent frontend of the Archive, reachable from https://materialscloud-archive-failover.cineca.it/.
- The new Archive backend has already been deployed at CINECA, in a development environment within the same cloud allocation of point 2.

For what concerns the CINECA site, all these deployments have been made using the cloud allocation described in Section 2, as depicted in Figure 5.

Regarding ongoing and future work, two main goals are being pursued: (a) finalize the CINECA failover service and (b) make the Archive data available to the users of the backup sites via an API. In particular, point (a) necessitates to:

(a.1) implement an automatic DNS switch in case of failure of the main, CSCS-hosted, frontend;

(a.2) switch from the old to the new Archive backend as soon as the main frontend does.

The same steps are likely to followed from the Jülich centre as well, so to deploy a second, Germany-based, failover of the Materials Cloud Archive. Regarding point (b), the switching to the new InvenioRDM backend (point a.2) directly provides us with an API to query the local database. A testing phase will follow in which the API will be restricted to read-only mode (as the Archive failover), and load tests will be performed to ensure stability of the overall infrastructure.

References

[1] Talirz, L. *et al.* Materials cloud, a platform for open computational science. *Scientific data* **7**, 299 (2020).

[2] Materials Cloud. https://www.materialscloud.org/home.

[3] Wilkinson, M. D. *et al.* The fair guiding principles for scientific data management and stewardship. *Scientific data* **3**, 1–9 (2016).

[4] Materials Cloud data management plan. https://www.materialscloud.org/dmp.

[5] Re3data. https://www.re3data.org/repository/r3d100012611.

[6] Eudat b2find. https://b2find.eudat.eu/organization/about/materialscloud.

[7] The Invenio digital library framework. https://invenio.readthedocs.io/en/latest/index.html.

[8] The inveniordm application. https://inveniordm.docs.cern.ch/.

[9] CINECA ADA Cloud documentation. https://wiki.u-gov.it/confluence/display/SCAIUS/ADA.

[10] rclone copy documentation. https://rclone.org/commands/rclone_copy/.

[11] JFZ JUDAC object storage. https://apps.fz-juelich.de/jsc/hps/just/object-storage.html.

[12] Ansible. https://www.redhat.com/en/ansible-collaborative.

[13] CINECA ADA Cloud security guidelines. https://wiki.u-gov.it/confluence/display/SCAIUS/Security+guidelines.

[14] Zenodo. https://zenodo.org/.

[15] Bastonero, L. *et al.* First-principles hubbard parameters with automated and reproducible workflows. *npj Comput. Mater.* **11** (2025).

[16] Uhrin, M., Zadoks, A., Binci, L., Marzari, N. & Timrov, I. Machine learning hubbard parameters with equivariant neural networks. *npj Comput. Mater.* **11** (2025).