

nffa.eu

PILOT 2021 2026

DELIVERABLE REPORT

WP9 Virtual Access

D9.2

Integration of the first VA services into the platform

Due date

M31



This initiative has received funding from the EU's H2020 framework program for research and innovation under grant agreement n. 101007417, NFFA-Europe Pilot Project

PROJECT DETAILS

PROJECT ACRONYM

NEP

PROJECT TITLE

Nanoscience Foundries and Fine Analysis - Europe|PILOT

GRANT AGREEMENT NO:

101007417

FUNDING SCHEME

RIA - Research and Innovation action

START DATE

01/03/2021

WORK PACKAGE DETAILS

WORK PACKAGE ID

WP9

WORK PACKAGE TITLE

Virtual Access

WORK PACKAGE LEADER

Dr. Rossella Aversa (KIT)

DELIVERABLE DETAILS

DELIVERABLE ID

D – D9.2

DELIVERABLE TITLE

Integration of the first VA services into the platform

DELIVERABLE DESCRIPTION

This deliverable presents the list of Virtual Access services currently integrated and provided in NEP.

DUE DATE

M31 30/09/2023

ACTUAL SUBMISSION DATE

DD/MM/YYYY

AUTHORS

R. Aversa (KIT), M. Panighel (CNR-IOM), A. Recchia (eXact lab), Y. Chiuconi (eXact lab), G.P. Brandino (eXact lab), G.D. Tsibidis (FORTH), Y. Pantazis (FORTH), I. Sfakianakis (FORTH), D. Grassano (EPFL)



PERSON RESPONSIBLE FOR THE DELIVERABLE

Dr. Rossella Aversa (KIT)

NATURE

- R - Report
- P - Prototype
- DEC - Websites, Patent filing, Press & media actions, Videos, etc
- O - Other

DISSEMINATION LEVEL

- P - Public
- PP - Restricted to other programme participants & EC: (Specify)
- RE - Restricted to a group (Specify)
- CO - Confidential, only for members of the consortium



REPORT DETAILS

ACTUAL SUBMISSION DATE

DD/MM/YYYY

NUMBER OF PAGES

13 (right-click and select "update the field")

FOR MORE INFO PLEASE CONTACT

Rossella Aversa,
 Karlsruhe Institute of Technology
 Hermann-von-Helmholtz-Platz 1
 76344 Eggenstein-Leopoldshafen
 Germany

email: rossella.aversa@kit.edu

VERSION	DATE	AUTHOR(S)	DESCRIPTION / REASON FOR MODIFICATION	STATUS
1	01/09/2023	R. Aversa	First draft	Draft
2	08/09/2023	R. Aversa, A. Recchia	Sections 1 and 2	Draft
3	15/09/2023	R. Aversa, G. Tsibidis, I. Sfakianakis, D. Grassano	Sections 3 and 4	Draft
4	20/09/2023	R. Aversa	Final revision	Finalised

CONTENTS

Introduction	5
1. MetaRepo	5
2. TriDAS (Trieste Advanced Data Services)	7
3. Modeling and Machine Learning	8
Machine Learning: Predictive service for nanostructures	8
Materials Modeling: damage threshold evaluation	9
4. Materials Cloud	11
References	13



INTRODUCTION

This deliverable presents the list of Virtual Access (VA) services currently provided in NEP. The offer comprises online simulation and machine learning services, as well as data and metadata services, integrated into the NEP infrastructure and provided as cloud services to authenticated users. The access is provided by the facilities hosting the services. Each Access Provider operates its own infrastructure, made of one or more installations, which offers identified User Access to the VA services under its own scheme based on institutional policies. Most of the VA services have been developed in the Joint Activity 6 on FAIR data approach. For these services, the deliverable in which they have been described is reported.

1. MetaRepo

MetaRepo [1] is a generic metadata repository and metadata schema registry, based on MetaStore [2], developed at KIT and running on a virtual machine (4 CPUs, 16 GB RAM, 50 GB SDD, OS: Debian 11, SSL via Apache reverse proxy) hosted at KIT.

It enables data curators to register metadata schemas in XML Schema Definition (XSD) or in JSON Schema, and it allows research users to store metadata documents which are automatically validated at upload time against the corresponding registered metadata schema. MetaRepo offers an Application Programming Interface (API), a software-to-software interface that enables other applications to exchange data with it, and a Graphical User Interface (GUI), suitable for users (Figure 1).

MetaStore Frontend for NFFA EU Pilot
 Schema and Metadata Management

Schema Management Metadata Management Search Show/Hide Filters Not logged in Login

Metadata Documents

Identifier	Related Resource
▼ mldata_basic_schema (version=1) (1 item)	
2b649aca-2d52-47f8-81d7-c8a1d9348bbe	https://b2share.eudat.eu/api/files/5fc88ad5-2f13-483c-8b80-a5862c91dbbb/Biological.tar
Schema Identifier mldata_basic_schema (version=1)	
Date Updated 2023-02-10 07:39	
▼ mri_schema (version=7) (2 items)	
3cbb28bd-62e7-425a-9024-288348f0741c	https://doi.org/10.5281/zenodo.7778338
Schema Identifier mri_schema (version=7)	
Date Updated 2023-03-28 15:06	
5bc69277-711c-4eb0-937d-d81c59dbbe36	https://doi.org/10.5281/zenodo.6107721
Schema Identifier mri_schema (version=7)	
Date Updated 2023-03-28 15:05	
▼ mldata_basic_schema (version=2) (2 items)	

Register new Metadata Document

Figure 1: MetaRepo homepage, where the login is available and metadata schemas or metadata documents can be browsed.

All publicly available resources can be seen from the frontend without logging in, while visualisation or editing of private metadata documents can be performed only by authenticated users, according to the Access Control List (ACL). Authentication is done via the NEP Single Sign-On system.

The usage of the service is monitored by aggregating the Units of Access (UoA) to the service. The UoA for the MetaRepo is defined as the single call or action (e.g. upload, query, retrieve) performed on the service by an authenticated user. MetaRepo uses a filter to detect each call to the service: each access detected by the filter is logged by the NEP backend, whereby as minimal private data as possible is stored. Currently, only the ID of the authenticated user and the ID of the VA service are transmitted. The NEP backend provides a REST endpoint to collect this data from an authenticated user, and additional data can also be provided.

As pointed out in the NEP privacy policy [3], the MetaRepo collects monitoring data only in order to send them to the NEP backend, without storing them locally. As a matter of fact, the MetaRepo stores ACLs for every metadata document, which contain the proposal ID and the user ID in order to grant/deny access to a specific resource. Any sensitive data contained in metadata documents stored in the MetaRepo are neither checked or analysed, and are under the responsibility of the creator of the resource.

The technical details about MetaRepo have been already described in Deliverable D9.1 [4], however some variation and additional features have been implemented:

- A generic search user interface has been implemented and can be extended according to the user feedback
- In addition to the already existing ones, a set of minimal metadata schemas has been registered for the description of Input, Precursor, Sample, Raw Data, Analysed Data. They include a basic, common list of attributes which will be extended according to the user feedback. These schemas also contain a provenance model, which include the reference to the internal ID of one or more parent metadata documents already stored in MetaRepo.
- The naming convention implementation has been modified after the received feedback: a human-readable label is assigned to the registered metadata schemas, to identify their type (Proposal, Input, Precursor, Sample, Raw Data, Analysed Data), while the metadata documents IDs are automatically assigned from MetaRepo (and no longer manually edited). The metadata documents are always associated with the corresponding metadata schema they are based on.
- A tool has been created to automatically dump the NEP proposal information, collect all the approved and finalised proposals (first trigger: after feasibility finalisation, second trigger: after the access has been concluded), map the information to the proposal metadata schema (registered in MetaRepo), create the metadata documents, and upload them to MetaRepo, assigning READ access to all the proposal members. A service account has been centrally created to perform the needed authenticated calls to MetaRepo. The proposals will be registered after the current testing phase.
- In addition to the MetaRepo GUI, a local Metadata editor [5] has been developed to allow users who are not familiar with JSON to load the metadata schemas registered in MetaRepo, to edit their metadata documents via a friendly interface and to register them in MetaRepo. Moreover, the creation of a metadata document containing the provenance workflow of a given resource can be automatically triggered. A view of the Metadata editor initial page is shown in Figure 2.



Easily edit metadata schemas

MetaStore schemas
Local documents
Provenance

Figure 2: Overview of the Metadata Editor initial page

2. TriDAS (Trieste Advanced Data Services)

TriDAS [6] is a set of data services based on data analytics and machine learning developed by CNR-IOM and eXact lab, hosted on the High Availability (HA) Kubernetes cluster of the ORFEO HPC Facility at AREA Science Park. It enables users to interactively explore scientific images using metadata, providing downloadable plots as well as content-based image retrieval. Additionally, it offers machine learning algorithms for the classification of SEM images.

TriDAS offers an intuitive user experience with three distinct types of sub-services, described in details in Deliverable D16.2 [7] and browsable from the homepage (see Figure 3):

- SEM and STM Explorers: These sub-services enable the users to explore datasets by selecting metadata pairs from drop-down menus, then interactive plots allow users to choose a subset of images for downloading.
- SEM and STM Similarity: users can select or drag and drop their own images, and the sub-service returns the most similar images from the database. Users have the option to select and download these images.
- SEM Classifier: users can either select or drag and drop user images, and the sub-service provides the category rate as results.

The service is authenticated upon the NEP Single Sign-On system: users can access the sub-services page without logging in, but they must authenticate to use the services. The usage is quantified by UoA, which are established to be every single action made by a logged-in user on one of the sub-services, and monitored: whenever a logged-in user performs an action on a sub-service, such as classifying a SEM image with SEM classifier or searching STM data using STM Explorer, the TriDAS backend sends a REST request to the NEP backend. This request includes both the service and sub-service ID and increases the usage counter by 1 UoA.

No information about the users are handled or stored by the service. The Keycloak token, used by the Single Sign-On, is the only piece of information needed to grant access to the service.

According to the proposal, TriDAS was planned to be fully operational at Month 31, however it was released in advance, at Month 24.

Trieste Advanced Data Services

Trieste Advance Data Services (TriDAS) is designed to enable advanced data services. It is implemented in Python and uses Bokeh for data visualization, the spym package to process and plot images, and the FastAPI framework as the backend. The platform host several virtual access services that can be classified in two broad categories:

- 1) Data Services developed within specific experimental/computational techniques;
- 2) Data Services are developed at a more general level to address a specific class of scientific data (i.e. images, 3D archives).

[Browse The Services](#)




Figure 3: TriDAS homepage, where the login is available and the services can be browsed.

3. Modeling and Machine Learning

FORTH offers to logged-in users two VA services: a predictive service for nanostructures and a damage threshold evaluation, described in detail in Deliverable D16.4 [8]. The services, authenticated upon the NEP Single Sign-On system, are accessible and can be browsed from a common GUI [9]. Their main features are illustrated in the following subsections.

The usage of each of the services is quantified by UoA, defined as one week (24 hours x 7 days = 168 hours) of computing power, and monitored: after 168 hours of service usage by logged-in users, the backend issues a REST call (including the ID of the service) to the NEP backend in order to increase the usage counter by 1 Unit of Access. FORTH maintains in the database a table for each of the services, storing the user data inserted at registration time (e.g., ID, first name, last name, email), plus an additional table which collects the user IDs and the incrementing hours of service usage. When 168 hours are reached (corresponding to 1 UoA), the service issues a REST call to the NEP backend and resets the counter.

Machine Learning: Predictive service for nanostructures

The Predictive Service, running on a virtual machine (2 CPUs, 4 GB RAM, 1024 GB SDD, OS: CentOS 7.9.2009) hosted at FORTH, allows to predict morphological nanostructure of laser-processed surfaces from Machine Learning (ML) models, which have been trained on annotated datasets. The predictive model of the service takes as input the natural logarithm of the number of pulses and the fluence, and returns the most probable surface structure. Feature engineering is performed, and new higher-order features are constructed and then fed to the ML models. The software trains several ML models including k-Nearest Neighbours, Gaussian Naive Bayes, Logistic Regression, Support Vector Classifier and Gradient Boosting Classifier.

There are two usage scenarios for the predictive service:

- Usage scenario 1: the use of trained ML models on new input data to predict their most probable surface pattern (Figure 4). The available predictive models have been trained on the following materials: Si, Steel 1.7131, Steel 1.7225 and Ti6Al4V. The expected input file (in csv format) should contain two columns corresponding to the input variables (natural logarithm of the number of pulses and fluence).




Figure 4: Usage scenario 1: prediction of the most probable surface pattern

- Usage scenario 2: the training of the ML models on new materials (Figure 5). The expected input file (in csv format) should contain three columns: natural logarithm of the number of pulses, fluence, and the label of the ground-truth surface pattern corresponding to the input variables. All the available ML models are then trained on the new input data and a final report, containing the mean accuracy and the corresponding standard deviation from a 5-fold cross-validation procedure, is generated. The best model is saved and can be used at later times for predictions (usage scenario 1) on the same material.




Figure 5: Usage scenario 2: training of the ML models on new materials

Materials Modeling: damage threshold evaluation

The service offers an interface (Figure 6) to perform in a user-friendly way the evaluation of the impact of various parameters such as the photon energies, the pulse duration, the pulse separation (in case of double pulse experiments) and the material thickness on the damage threshold for various metals (Au, Ag, Cu, Al, Ni, Ti, Cr, Stainless Steel); three different substrates (Si, SiO₂ and soda lime silica glass) have been considered.

Material	Substrate	Thickness (nm)	Lambda (nm)	TP (fs)	Delay (ps)
Nickel	SiO2	1030	10	170	0.020

Run

Figure 6: Interface for evaluation of damage threshold.

A multiscale physical model is used that correlates the energy absorption, electron excitation, relaxation processes and minimal surface modification. The model takes into account both the thermophysical and optical parameter values of the irradiated complex (metal/substrate). It was developed assuming the most common metals and substrates, and it is aimed to be extended to other metals and configurations (multi-layered materials, etc).

Results are saved in txt format files (Figure 7) that contain the prescribed laser conditions and the materials used as well as information of the damage threshold (i.e. minimum fluence to melt the material). The maximum temperature in these conditions and comparison with the melting temperature is used to assess the accuracy of the algorithm. Furthermore, optical parameters (i.e. reflectivity and absorptivity) of the irradiated metals are calculated and their values are included in the txt file. A step-by-step documentation [10] is available in the interface, together with the details of the underlying theory that describes the multiscale physical processes that lead to materials damage.

Results:

1	Material:	Ni
2	Substrate:	SiO2
3	Wavelength (nm):	1030
4	Thickness (nm):	10
5	Damage Threshold (J/cm2):	0.035635
6	Discretization size (nm):	0.43333
7	Pulse duration (fs):	170
8	Pulse separation (fs):	20000
9	Maximum Temperature(K):	1728.0008
10	Damage point (K):	1728
11	Reflectivity:	0.33268
12	Absorptivity:	0.37088
13	Melting point (K):	1728

Figure 7: Output of the simulation for the damage threshold evaluation

The software performing the simulations is written in Octave and the code is executed on a computer (AMD Ryzen Threadripper 1950X 16-Core Processor 3.40 GHz, 64.0 GB) located at FORTH.

4. Materials Cloud

Materials Cloud is a platform designed to enable open and seamless sharing of resources for computational science, driven by applications in materials modelling. It hosts AiiDALab [11], an open-source web platform that provides a Jupyter-based environment pre-configured with the AiiDA materials informatics infrastructure [12] for automated workflows and provenance tracking, as well as a number of other tools that are commonly used in computational materials science.

The service, running on a virtual machine (32 CPUs, 128 GB RAM, 1 TB storage, OS: Ubuntu-20.04, SSL via Apache reverse proxy) hosted at CSCS, is provided to NEP users with a preinstalled version of the Quantum ESPRESSO code [13] and the corresponding app [14], which enable users to run density-functional theory calculations using turn-key workflows and to compute properties of materials, such as the electronic band structure and the density of states and equilibrium geometries.

Users are able to access the service and authenticate by connecting to the web GUI [15]. Figure 8 shows the initial page of the service after login. The GUI and the workflows are designed to be user friendly and accessible, allowing users with limited knowledge on simulations to compute material structural and electronic properties given an input structure. Through the GUI, users can visualise the results of the calculations, as well as download the results, the raw data and the provenance information.

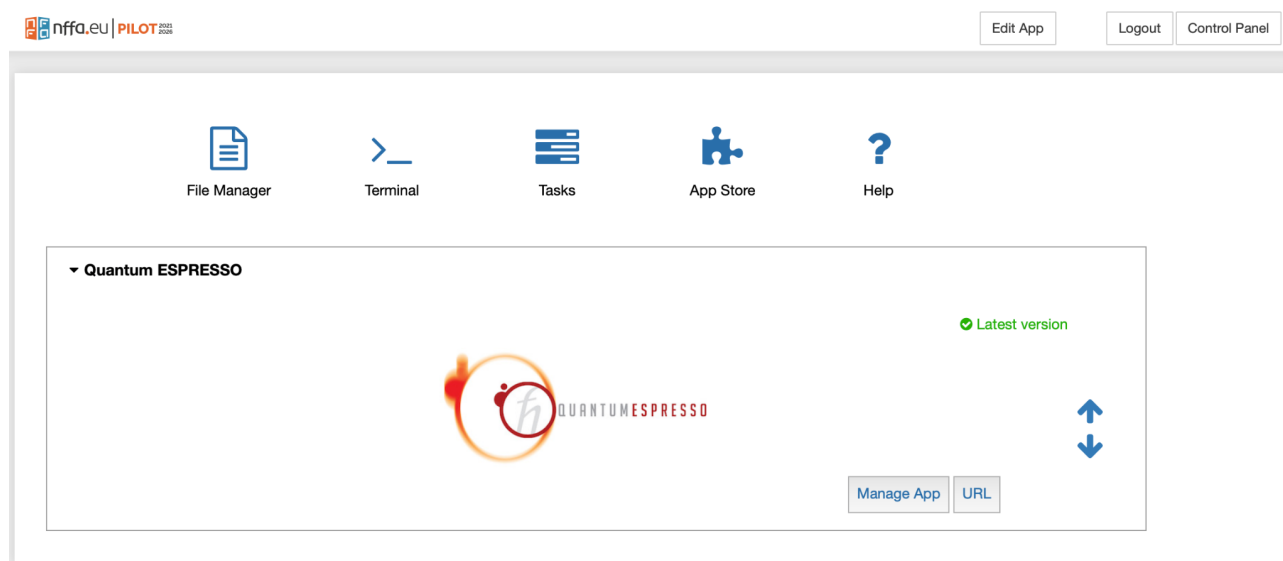


Figure 8: The initial page of Materials Cloud AiiDALab after authentication

From the preinstalled app, users are able to select a structure (Figure 9), either by uploading a structure file (a wide variety of commonly used formats is supported [16]) or by directly querying a materials database, such as the Materials Cloud or the Materials Project, via the OPTIMADE interface.

Figure 9: AiiDALab Quantum Espresso app interface for selecting an input structure

After defining the input structure, the user selects the workflow to run, with the desired properties to be computed, and the resources on which the calculation will be performed. These can vary from the virtual machine itself for small calculations (limited to 4 cores and 12 GB RAM per user), to SSH accessible HPC resources, that the user will need to link to the service (instructions on how to proceed are available both in the GUI and at [17]).

Authentication of the users takes place using the NEP Single Sign-On, and the usage will be monitored following the agreed definition of UoA as the number of authenticated users accessing the service. Specifically, whenever a user logs into the service (spawning or connecting to an existing deployed container), the server of the service sends a HTTP call to the NEP backend increasing the counter by 1 UoA. The monitoring data is collected only for the purpose of sending it to the NEP backend monitoring server without being stored locally in order to comply with the NEP privacy policy [3]. The data generated by the users is stored on the virtual machine and identified using an UUID only.

References

- [1] MetaRepo homepage: <https://metarepo.nffa.eu>
- [2] MetaStore documentation: <https://github.com/kit-data-manager/metastore2/blob/metastoreGui/restDocu.md>
- [3] NEP privacy policy: <https://nffa.eu/apply/privacy/>
- [4] Aversa, R., Brandino, G. P., Chelbi, S., Hartmann, V., Jejkal, T., NFFA-Europe Pilot - D9.1 - Deployment of the VA service prototypes. DOI: 10.5281/zenodo.7702265
- [5] Metadata Editor GitLab page: <https://metadata-editor.gitlab.io/documentation>
- [6] TriDAS homepage: <https://tridas.nffa.eu>
- [7] Franzon, M., Osmenaj, E., Panighel, M., Rodani, T., Brandino, G. P., Aversa, R., NFFA-Europe Pilot - D16.2 - Report on the first data services. DOI: 10.5281/zenodo.7057941
- [8] Aversa, R., Blumenröhr, N., Recchia, A., Tsididis, G. D., Pantazis, Y., NFFA-Europe Pilot - D16.4 - Report on additional data services. Submitted
- [9] FORTH services homepage: <https://nffa-modeling.iesl.forth.gr/>
- [10] Tsididis, G. D., Manual for the use of the interface to compute the damage threshold of metals of various thicknesses following irradiation with femtosecond laser pulses. <https://nffa-modeling.iesl.forth.gr/docthermal>
- [11] AiiDALab website: <https://www.aiidalab.net/>
- [12] AiiDA website: <https://www.aiida.net/>
- [13] Quantum ESPRESSO homepage: <https://www.quantum-espresso.org/>
- [14] AiiDALab Quantum ESPRESSO app: <https://github.com/aiidalab/aiidalab-qe>
- [15] Materials Cloud AiiDALab homepage: <https://aiidalab-nep.materialscloud.org>
- [16] AiiDA supported formats: <https://wiki.fysik.dtu.dk/ase/ase/io/io.html#ase.io.write>
- [17] AiiDALab documentation <https://aiidalab-qe.readthedocs.io/?badge=latest>

