



Scientific Visualization on the Cloud: the NEANIAS Services towards EOSC Integration

Eva Sciacca · Mel Krokos · Cristobal Bordiu · Carlos Brandt · Fabio Vitello · Filomena Bufano · Ugo Becciani · Mario Raciti · Giuseppe Tudisco · Simone Riggi · Eugenio Topa · Sami Azzi · Benjamin Kyd · Simone Mantovani · Laura Vettorello · Jiacheng Tan · Josep Quintana · Ricard Campos · Noela Pina

Received: 7 July 2021 / Accepted: 11 January 2022
© The Author(s) 2022, corrected publication 2022

Abstract NEANIAS is a research and innovation action project funded by the European Union under the Horizon 2020 program. The project addresses the challenge of prototyping novel solutions for the underwater, atmospheric and space research communities, creating a collaborative research ecosystem, and contributing to the effective materialization of the European Open Science Cloud (EOSC). NEANIAS drives the co-design, implementation, delivery, and integration into EOSC of innovative thematic and core services, derived from state-of-the-art assets and practices in the target scientific communities.

The research leading to these results has received funding from the European Commission's Horizon 2020 research and innovation programme under the grant agreement No. 863448 (NEANIAS).

E. Sciacca (✉) · C. Bordiu · F. Bufano · U. Becciani · M. Raciti · G. Tudisco · S. Riggi
INAF, Catania Astrophysical Observatory, Via S Sofia 78,
95123 Catania, Italy
e-mail: eva.sciacca@inaf.it

F. Vitello
INAF, Institute of Radioastronomy, Via Gobetti 101,
40129 Bologna, Italy

M. Krokos · B. Kyd · S. Azzi
School of Creative Technologies, University of Portsmouth,
Eldon Building, Portsmouth PO1 2DJ, UK

C. Brandt
Jacobs University Bremen, gGmbH Campus Ring 1,
28759 Bremen, Germany

We present the overall NEANIAS ecosystem architecture, with an emphasis on its core visualization services, detailing their specifications and software development plan, and focusing on the underpinning service-oriented architecture for their delivery. We report on the underlying ideas and guiding principles for designing such visualization services, outlining their current release status and future development roadmaps towards Technological Readiness Level (TRL) 8 maturity and EOSC integration.

Keywords Scientific visualization · Cloud services · Service oriented architectures · Big data · European Open Science Cloud (EOSC)

E. Topa
ALTEC Spa, C.so Marche 79, 10146 Torino, Italy

J. Tan
School of Computing, University of Portsmouth,
Buckingham Building, Portsmouth PO1 3HE, UK

S. Mantovani · L. Vettorello
MEEO S.r.l., Corso Ercole I d'Este 6/A - 44121
Ferrara, Italy

J. Quintana · R. Campos
CORONIS Computing S.L., Carrer Pic de Peguera 11,
17003 Girona, Spain

N. Pina
Ubiwhere Lda, Coimbra, Portugal

1 Introduction

The European Cloud initiative¹ aims at providing science, industry and public authorities with world-class infrastructures for big data storing, managing, transporting and processing, in the expectation of significantly streamlining exploitation of open science practices.² The initiative is built around two main pillars: the European Open Science Cloud (EOSC)³ which is a virtual environment for handling the large volumes of information generated by the big data revolution; and the European Data Infrastructure (EDI),⁴ which underpins EOSC through the deployment of high-bandwidth networks and supercomputing resources, enabling the access and processing of big data volumes in the Cloud.

The EOSC roadmap⁵ foresees the seamless federation of existing and future research data infrastructures, as well as other large-scale scientific networks and initiatives. Its vision envisages the adoption of a common policy framework to enable FAIR⁶ (i.e. Findable, Accessible, Interoperable and Reusable) data to be used and shared throughout the entire value chain for scientific, societal and industrial purposes. The NEANIAS⁷ project emerges in this context as an ambitious research and innovation action specifically addressing the prototyping of new innovative services⁸ for underwater, atmospheric and space research thematic sectors. Starting from TRL 6 software solutions, the project is developing, validating, delivering, and maturing into TRL 8 a portfolio of cutting-edge thematic services⁹ for EOSC onboarding to fill current research gaps and needs, engaging user communities and identifying novel business cases and opportunities. Although there are evident differences among the NEANIAS scientific sectors, its thematic services

are built upon a comprehensive collection of reusable *core* services, designed to tackle fundamental technical or market-oriented data life cycle tasks. Due to their agnostic nature, utilization of such core services extends well beyond the initial scope of NEANIAS for exploitation by other research communities in academia and industry, as well as public authority bodies. This article describes the overall architecture, specifications and software development plan of NEANIAS, describing thematic and core services and release status focusing on visualization.

The paper organisation is as follows. Section 2 outlines the main challenges of scientific visualization in the underwater, atmospheric and space research sectors. Section 3 refers to the related works. The overall architecture of the NEANIAS ecosystem is described in Section 4, where we discuss concepts relating to the service-oriented architecture, REST¹⁰ (REpresentational State Transfer) paradigm, standards and interoperability, as well as logical architecture and generic service processing lifecycles. Section 5 focuses on core visualization services, summarising the technologies adopted and including possible dependencies with other core services. The service delivery model is outlined in Section 6 discussing also services evaluation. Section 7 outlines the current release status. We conclude with a summary of the work carried out so far and brief insights into future developments.

2 Scientific Visualization Challenges

Visualization plays a fundamental role in almost every scientific discipline facilitating qualitative and quantitative data analysis, for new knowledge generation and effective communication of end results. Suitable tools and approaches can boost scientific productivity significantly, e.g., by revealing hidden trends or intrinsic patterns in the data, leading to fresh insights and eventually, new scientific discoveries.

The big data revolution is providing enormously large, incredibly rich, and highly complex data volumes that impose extremely challenging demands on traditional visualization approaches (see e.g. [9, 14]). The demands to address are *efficiency*, i.e. the ability to handle rapidly the underlying data complexity,

¹<https://ec.europa.eu/digital-single-market/en/cloud>

²<https://www.openaire.eu>

³<https://www.eosc-portal.eu>

⁴<https://eudat.eu>

⁵https://ec.europa.eu/research/openscience/pdf/eosc_strategic_implementation_roadmap_short.pdf

⁶https://ec.europa.eu/info/sites/info/files/turning_fair_into_reality_0.pdf

⁷<https://www.neanias.eu/>

⁸<https://eosc-portal.eu/infraeosc-02-2019-prototyping-new-innovative-services>

⁹<https://catalogue.neanias.eu>

¹⁰<https://restfulapi.net>

and *intuition*, i.e. the ability to reach suitable interpretation by domain experts. The cluster of core visualization services being developed in NEANIAS (see Section 5) is founded firmly on these guiding principles to underpin a portfolio of thematic services¹¹ which are tailored to specific community challenges in the underwater, atmospheric and space research sectors, as outlined below.

2.1 Underwater Services Technical Challenges

The underwater community is multidisciplinary, encompassing among others archaeologists, environmental scientists, marine geologists, industrial energy planners and engineers. The challenge is to deploy appropriately topographical and geological information of the seafloor for a variety of purposes e.g. as diverse as marine habitat mapping, biodiversity studies, tectonic plate analysis and fault detection, finding and protecting archaeological sites or even oil pipeline planning. The data context is underpinned by the rapid growth of underwater mapping sensors which enables retrieving the shape of the Earth's seafloor at unprecedented levels of fidelity. Computationally-efficient visualization of such datasets is needed for underwater communities, by generating world-scale multi-resolution triangulated irregular networks optimized for web-based visualization [6].

The idea is to transmit the data and defer visualization on the user's end, where the graphical hardware of the client can be used focusing on striking a balance between the amount of data transfers and efforts required for rendering. Then, level of detail techniques, able to change the complexity of the displayed data based on the point of view desired by the end user, are suitable for deployment to focus on rendering the part of the world falling in the user's frustum with a complexity that adapts to the distance from viewers. Although bathymetry can be visualized as raster 2D maps, it provides information about the 3D relief of the seafloor, and as such it is better interpreted via rich 3D visualizations. Transferring and visualizing 3D data however is more involved since instead of images triangle meshes are to be rendered. To cope with these challenges, the NEANIAS visualization core services

provide specific tools for generating large hierarchical tiled data structures for imaging and bathymetric maps (see Section 5.1.3). We adapt state-of-the-art simplification approaches to efficiently work at tile level, for highly parallelizable rendering paying attention on ensuring 3D tile continuity in their hierarchy within the same level of detail.

2.2 Atmospheric Services Technical Challenges

The atmospheric community includes a multitude of stakeholders e.g., urban air quality authorities that monitor the concentration of gases and small particles in the air, geologists who track gas fluxes from active faults or volcanoes and meteorologists who monitor the evolution of physical variables to provide weather forecasts. Such stakeholders are facing diverse challenges, e.g. in visualising and forecasting spatial distribution of pollution concentrations within urban areas. Typically, such concentrations are monitored through local air quality stations, however, the in-situ measurements can only give information of pollutants for specific areas of measurements. Air quality modelling methodologies are then essential, albeit very challenging as concentration levels of atmospheric pollutants arise as a result of local contributions (i.e. from specific emission sources) and include additional background pollutant concentrations transported from surrounding regions. The spatial distribution of the emissions induced by local sources, especially road traffic, also depends on the wind conditions due to buildings and urban canyon structures.

The challenge is then to define spatial resolution modelling approaches striking an appropriate balance between the availability of input data with high detail and the ability to visualize pollutant concentrations at a spatial resolution sufficient for urban air quality planning [7]. The underlying modelling system needs to take into account local emissions modelling, weather predictions and urban air quality modelling that includes local scale effects and transport of pollutants from long distances. To address such challenges, NEANIAS provides to the atmospheric community cloud services for high resolution mapping of pollutant concentrations on specific areas (see Section 5.4). These services also provide weather forecasts, enabling to visualise within pollution maps several relevant parameters, e.g. temperatures and precipitations.

¹¹<https://docs.neanias.eu/en/latest/>

2.3 Space Services Technical Challenges

The space community includes astrophysicists and planetary scientists, who are facing a data explosion revolution fuelled by an increasing number of planetary exploration missions (e.g., the Perseverance¹² rover on Mars) and the upcoming deployment of new state of the art observing facilities (e.g., the Square Kilometre Array - SKA,¹³ the Vera C. Rubin Observatory,¹⁴ or the James Webb Space Telescope - JWST¹⁵). These instruments are expected to produce overwhelming amounts of data volumes, by mapping solar system bodies with an unprecedented detail, or performing new all-sky surveys with unparalleled sensitivity and angular resolution [20]. Visual exploration and discovery mechanisms can be instrumental in dealing efficiently with this avalanche of large-scale datasets expected to be delivered from new instrumentation and missions.

The main technical challenges for the space community are centered around data accessibility, interoperability and scalability. While the use of standard data formats is accepted widely, e.g. FITS [24], observatories usually maintain their own archive interfaces, relying on custom made (i.e. not always standardized) search keywords. Such an heterogeneity constitutes an important obstacle that may slow down the generation of scientific results, especially for multiwavelength astronomy –the study of astronomical objects combining information acquired at various wavelengths– that requires simultaneous visualization of images typically obtained with different instruments (and in very different conditions). The International Virtual Observatory Alliance (IVOA¹⁶) has made significant progress in this direction towards a true standardization of astronomical data access, providing query interfaces for multiple catalogues and data archives, and developing protocols for data access and exchange (e.g. the Table Access Protocol, or TAP¹⁷). However, challenges remain, e.g. use of IVOA tools is still far from being a preferred option as evidenced in a recent

survey of the European astrophysics community [4]. The demand for scalability arises from the continuously increasing size of datasets as efficient visualization mechanisms are impractical with standard hardware, unless novel rendering strategies exploiting next generation infrastructures are designed and deployed. Recent developments, e.g. CARTA,¹⁸ rely on client-server architectures, in an attempt to alleviate this performance bottleneck, paving the ground for successfully handling large data volumes from SKA and its pathfinders.

To cope with these challenges, the NEANIAS visualization core services provide tailored and scalable visual discovery workflows to be exploited by the Space Thematic Services, in order to meet the demands of the astrophysics community for advanced visual analytic solutions [4]. Additionally, the Open Science, FAIR-compliant approach followed is especially convenient in underpinning the standardization and interoperability of astronomical archives and databases. As an example, the ViaLactea Visual Analytics (VLVA¹⁹) service offers a complete visualization solution with seamless access to multiple surveys of the Milky Way (near-infrared to radio band) including continuum maps and molecular data cubes, that can be handled entirely by our Visual Discovery framework (see Section 5.1).

3 Related Works

Despite the intrinsic differences among the underwater, atmospheric and space research sectors, a series of common requirements arise. Innovative visualization tools and solutions must be able to (i) handle complex and heterogeneous datasets, (ii) support multiple visualization strategies (e.g., 2D and 3D renderings, projection techniques for higher dimensionality data) and (iii) enable an intuitive and user-friendly data exploration. Besides, the ever-growing size of the datasets underlines the need for moving from the traditional standalone model to novel distributed approaches, relying on cloud-based infrastructures able to meet the increasing demand for resources. In this sense,

¹²<https://mars.nasa.gov/mars2020/>

¹³<https://www.skatelescope.org>

¹⁴<https://www.lsst.org/>

¹⁵<https://www.jwst.nasa.gov/>

¹⁶<https://ivoa.net>

¹⁷<https://www.ivoa.net/documents/TAP/>

¹⁸<https://cartavis.org/>

¹⁹<https://marketplace.eosc-portal.eu/services/space-vis-vialactea-service>

we have recently witnessed the blossoming of several solutions and frameworks that are steadily gaining prominence in diverse scientific visualization fields.

ParaView [1] is a large scale parallel visualization software, designed for effective exploitation of high performance infrastructures. A web enabled version ParaViewWeb,²⁰ can act as a Web Application by allowing users to remotely connect via web browsers to a ParaView server. The connection is enabled via the custom library wslink²¹ that connects JavaScript web clients to a Python web server through ALP WebSockets. Furthermore, the in-situ library Paraview Catalyst [2] allows users to instrument their application for in-situ analysis, visualization, and computational steering purposes.

The Cactus computational framework [10] can support a web browser interface for insitu visualization and steering tasks. The user can instrument existing high performance applications with the Cactus API, perform steering tasks and view visualization outputs through a web browser. The standard implementation utilises an HTTPD²² web server, and forwards ports to the user for remote access.

WebVis [27] is a multi-user, client-server, visualization framework with a web-based client offering services in the cloud and is accessible via netbooks, smartphones, and other web-and JavaScript-enabled mobile devices. The client connects to a server via a back-end service built using the Google Web Toolkit²³ and a Java web server, which communicate via a web service that forwards events and images to and from the underlying rendering infrastructure nodes. Client GUI interactions are forwarded to the server, and images returned, through an EventBus using the HTTP server push paradigm.

Remote frameworks that support web such as FastX²⁴ are seeing some success in real-world deployments, e.g. for the web visualization portal at the Texas Advanced Computing Centre²⁵ which supports web-based usage of Paraview. FastX offers a cloud-capable display solution that is fully optimized in

terms of achieving high performances over high-latency connections. VisIVO Science Gateway [3] is wrapped around WS-PGRADE/gUSE [11] integrating services for processing and visualizing large-scale multi-dimensional astrophysical datasets on Distributed Computing Infrastructures.

Other cloud-based visualisation solutions are typically built upon bespoke frameworks which can be effective e.g. in the context of targeting single highly focused applications but less useful in providing a rich general purpose visualisation ecosystem. Core visualization services of NEANIAS offer such a visual discovery platform based on JupyterHub²⁶ allowing users to write and execute code, to transform, analyse, and visualize datasets by supporting several high performance visualisation services integrating seamlessly with cloud based mechanisms for authorisation and authentication, storage allocation and sharing, logging and accounting and open publication of resulting data products, thus making it ready for onboarding the emerging EOSC ecosystem.

Many of the aforementioned solutions certainly represent a significant step forward for scientific visualization, being widely adopted in several fields. Still, they constitute a somewhat dispersed, heterogeneous toolbox rather than a unified solution. Aspects such as data access and sharing, or interoperability and integration with other services and tools, are out of the scope of existing visualization suites. NEANIAS, with its visualization core services, intends to fill this gap, by providing a set of easy-to-use, multipurpose visualization frameworks built upon FAIR principles. NEANIAS core visualization services represent a state of the art solution for the visualization challenges of the big data era, facilitating straightforward data access and sharing by taking advantage of the EOSC cloud ecosystem and resources. Moreover, the core visualization services are not just independent pieces of software, but can also be seamlessly integrated with other tools by means of standard interfaces. This ability allows for building complex visualization analysis workflows from scratch. Last but not least, core visualization services, thanks to their domain-agnostic nature, can be easily exploited for purposes far beyond their original scope, i.e., in domains not related to underwater, atmospheric or space research, such as healthcare or communications.

²⁰ParaViewWeb:<https://kitware.github.io/paraviewweb/>

²¹wslink library:<https://github.com/kitware/wslink>

²²<https://httpd.apache.org/>

²³<http://www.gwtproject.org>

²⁴<https://www.starnet.com/fastx/cloud>

²⁵<https://www.tacc.utexas.edu/vislab>

²⁶<https://jupyter.org/hub>

4 The Neanias Ecosystem

The NEANIAS thematic services are being developed as totally independent software components to cater for specific community needs. Nevertheless, there is an evident number of commonalities that connect them intrinsically, e.g., when considering end-user workflows. We are exploiting such underlying synergies to establish generic, cross-community core services to amplify the potential of thematic services, pushing boundaries towards fully blown open science practices [21] and facilitating migration to EOSC concepts [19]. The NEANIAS core services are being developed under the umbrella of distinct and complementary clusters of specific functionalities for *open science lifecycles* (C1 services²⁷), *EOSC integration* (C2 services²⁸), *artificial intelligence* (C3 services²⁹) and finally, *visualization* (C4 services³⁰). Abstracting from thematic services underpinning concepts, the C3 and C4 clusters constitute the arrowhead of NEANIAS generic service offerings. C4 services in particular provide support mechanisms for high performance multidimensional data visualization solutions applicable to a raft of diverse scientific scenarios (see Section 5).

The design principles of the NEANIAS ecosystem are centred around the pillars of a) delivering a service-oriented architecture, while b) fully realizing the REST paradigm [8], and c) adhering to FAIR principles [25] to manifest data compliant with findable, accessible, interoperable and reusable mechanisms. The service-oriented architecture is realized through the use of microservices to enforce single-responsibility principles, allowing for component separation, reuse, discoverability, and interoperability. REST design patterns for application programming interfaces are well-established web standards commonly employed in existing EOSC services to satisfy architectural constraints such as client-server

stateless interactions, cacheable resources, layered systems and uniform interfaces. When applied to web services, these offer desirable properties, such as performance, scalability, and modifiability, that enable services to work on the Web optimally. We also support FAIR principles through the employment of technologies such as OpenAIRE's Zenodo platform³¹ and the eInfraCentral³² service catalogue, as well as adhering to standards and guidelines from international community bodies, e.g. the Open Geospatial Consortium (OGC)³³ and International Virtual Observatory Alliance (IVOA).³⁴

The overall architecture of the NEANIAS ecosystem follows a distributed and tiered approach (see Fig. 1). The model is composed of logical architecture blocks that cover specific functional or technological requirements and are appropriately exposed to other services (offered by NEANIAS or third-parties). Services themselves are also composite multi-tier systems encapsulating loosely coupled microservices and lightweight communication protocols for their internal operations. They provide appropriate abstractions of physical resources regarding computing, storage, and network capacity, e.g. virtual machines, container platforms, or storage programming interfaces. The NEANIAS thematic services encompass top-level offerings underpinned by core services to cater for specific community needs in the underwater, atmospheric, space and business sectors (tailored to energy and smart air quality) building use cases on top of thematic services. Core services include base and advanced offerings as follows:

- **Base Core Services:** This cluster of generic services provides base tooling in the ecosystem to support open science lifecycles (C1) and integration with EOSC and cloud resources (C2).

C1 services provide the tools for NEANIAS services to be discoverable, accessible and integrated with EOSC. The NEANIAS service catalogue portal and programming interfaces are based on the eInfraCentral catalogue and enable

²⁷NEANIAS C1 Services documentation: <https://docs.neanias.eu/en/latest/c1-services.html#c1-services>

²⁸NEANIAS C2 Services documentation: <https://docs.neanias.eu/en/latest/c2-services.html#c2-services>

²⁹NEANIAS C3 Services documentation: <https://docs.neanias.eu/en/latest/c3-services.html#c3-services>

³⁰NEANIAS C4 Services documentation: <https://docs.neanias.eu/projects/c4-services/en/latest/>

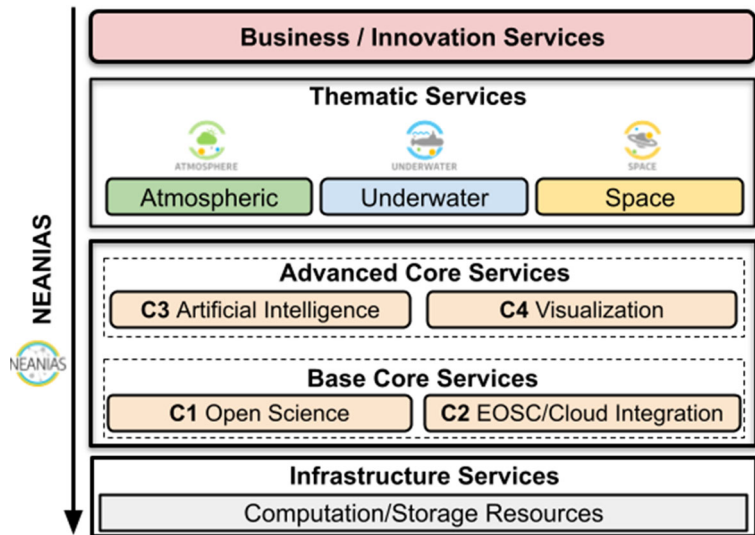
³¹<https://zenodo.org/>

³²<https://www.einfracentral.eu/>

³³<https://www.ogc.org/>

³⁴<http://ivoa.net/>

Fig. 1 Overall architecture of the NEANIAS ecosystem



service providers to register and present their services in a single project's catalogue, which is compliant with EOSC guidelines. The NEANIAS data catalogue is built on top of Zenodo and enables providers to publish their data as well as assign persistent identifiers (PID) for digital assets, making them unambiguously cited and hence univocally discoverable. Service providers are also able to access common UI and UX elements for the presentation of their services and development of uniformly coordinated user experiences for all services developed in the context of NEANIAS. Finally, thematic service providers can also create and maintain actionable data management plans through Argos³⁵ OpenAIRE service.

C2 services encompass loosely coupled elements that enable the assembly of virtual infrastructures to serve other services, bringing them closer to EOSC resources and processes, as well as other research and cloud computing infrastructures. The authentication and authorisation infrastructure³⁶ (AAI) service offers a horizontal solution for authenticated access (across users and services), a centralized user management, and an authorization policy definition, as well as an identity federation. It is backed by the open source identity and access management solution

Keycloak.³⁷ The configuration management³⁸ service provides key value storage for storing service configurations, allowing for integration at bootstrap and runtime. The instance registry³⁹ service supplies dynamic registration and service discovery, along with service location and health status information. The log aggregation⁴⁰ and accounting⁴¹ services provide seamless logging and accounting functionalities by taking advantage of the ELK stack.⁴² Finally, cloud computing services (supporting both Virtual Machines through OpenStack⁴³ and containers employing Kubernetes⁴⁴) and data services for data depositing, sharing and exploration are available.

- **Advanced Core Services:** This cluster is built upon base services for generic functionalities to be exploited by the NEANIAS specific domains and beyond; these include support for artificial

³⁷<https://www.keycloak.org/>

³⁸<https://docs.neanias.eu/projects/configuration-service/en/latest/README.html>

³⁹<https://docs.neanias.eu/projects/instance-registry-service/en/latest/README.html>

⁴⁰<https://docs.neanias.eu/projects/logging-service/en/latest/README.html>

⁴¹<https://docs.neanias.eu/projects/accounting-service/en/latest/README.html>

⁴²<https://www.elastic.co/what-is/elk-stack>

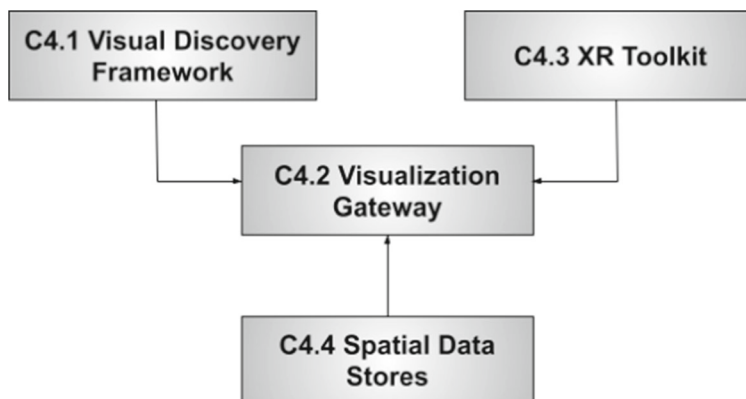
⁴³<https://www.openstack.org/>

⁴⁴<https://kubernetes.io>

³⁵<https://argos.openaire.eu/splash/>

³⁶<https://docs.neanias.eu/projects/aai-service/en/latest/>

Fig. 2 Overview of NEANIAS C4 visualization services



intelligence processing (C3) and advanced visualization solutions (C4). C3 services represent the elements of typical machine learning workflow lifecycles, initially a model needs to be designed and implemented. Whenever data loads exceed the capabilities of single workstations, model training is offloaded to distributed computational clusters. Depending on the model, different parallelisation approaches can be followed and once a model is trained, serving and deployment requires dedicated resources. Finally, the details of C4 visualisation services currently delivered within NEANIAS are discussed in the next section.

5 Core Visualization Services

The core visualization services in NEANIAS lie within the upper layer of the core services cluster in the tiered system architecture described in Section 4. They provide functionalities supporting different scenarios when integrated with the higher-level thematic services. Each service can be viewed as a fundamental building block of a typical scientific visualization workflow; the focus is on visual discovery frameworks (C4.1, VD - Section 5.1), science gateways (C4.2, VG - Section 5.2), Cross-Reality (XR) frameworks to support virtual reality (C4.3, XR - Section 5.3), and spatial data stores (C4.4, DS - Section 5.4).

The overall relationships and dependencies between the services are depicted in the schema in Fig. 2. The science gateway⁴⁵ enables visualization workflows based on the visual discovery framework and

accesses to the mechanisms exposed through the XR toolkit and the spatial data store services.

C4 services identify underlying common needs and synergies across user communities and apply cross-cutting glue and abstraction mechanisms to realize a robust, modular, cross-sector and open source core pool of methods and tools to support visualization services, leveraging on the resources available from research infrastructures and EOSC-hub⁴⁶ facilities. C4 services underpin a diverse range of multi-faceted visualization workflows, from 2D/3D spatio-temporal to composite 2D/3D visuals for complex, high dimensionality datasets, including computationally demanding cross-reality applications, e.g. supporting mechanisms for scalable visualization or production of assets to enable advanced interaction mechanisms through game engines and VR headsets.

Handling of big data relies on offloading workflow computations to distributed computational clusters. Depending on the tools deployed, different parallelization approaches can be exploited to harvest optimally different types of underlying infrastructures, ranging from small-scale clusters (with or without GPUs) to large-scale heterogeneous systems. Although the core services developed in C4 allow customization to meet NEANIAS specific community needs, they can be generalized sufficiently to extend to other EOSC end-user communities far beyond their original scoping.

5.1 C4.1 - Framework for Visual Discovery (VD)

C4.1 provides tools to support data-intensive computing for visual scientific discovery, including research,

⁴⁵<http://www.icsciencegateways.org/>

⁴⁶<https://www.eosc-hub.eu>

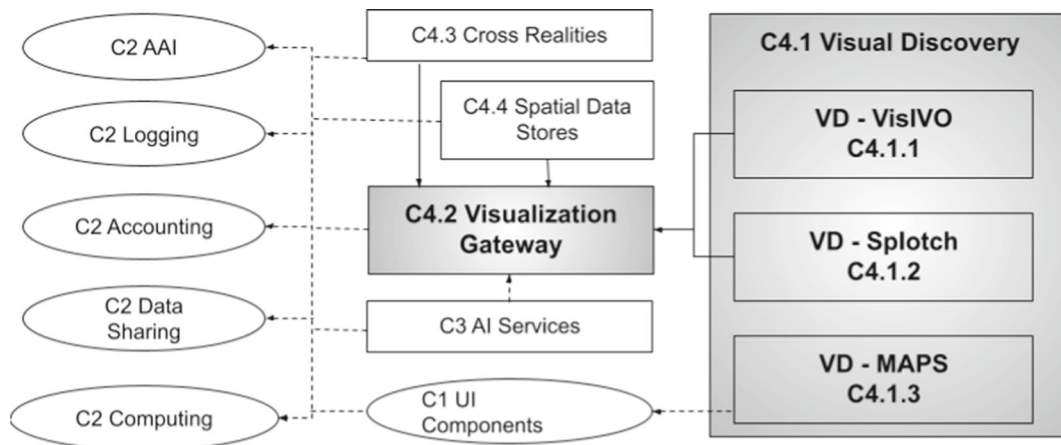


Fig. 3 NEANIAS C4.1 Framework for Visual Discovery services and C4.2 Visualization Gateway diagram and integration with Base Core services (C1 and C2) and C3 AI Services

training and outreach, on observational data, theoretical simulations, and 2D/3D tiles and maps. High-performance services are provided for: (1) visual analytics from multidimensional data tables (VD-VisIVO) (see Section 5.1.1); (2) high-quality volume rendering of particle-based datasets exploiting a variety of parallel programming models (VD-Splotch) (see Section 5.1.2); and (3) creation of interactive 2D/3D tilings and maps (VD-Maps) (see Section 5.1.3). Figure 3 shows the interdependencies and relationships with the base core services in the NEANIAS ecosystem. The first two visual discovery services VD-VisIVO and VD-Splotch are exposed to the end-users via the Visualization Gateway while the third one, VD-Maps is accessible through a specific web UI based on C1 common UI components. All VD services are integrated with C2 services for AAI, logging, accounting, data sharing and computing access. Particular details are summarized in the below sections.

5.1.1 VD-VisIVO

The VD-VisIVO service offers a framework for data-intensive visual discovery and analysis through the suite of tools provided by VisIVO [18]. VisIVO is designed specifically for distributed computing environments (e.g. Cloud infrastructures) and provides mechanisms for exploration of large-scale datasets through highly customized 3D views of multidimensional data tables from various sources.

5.1.2 VD-Splotch

The VD-Splotch service makes available the suite of tools provided by Splotch [15] for data processing and visual discovery. Splotch supports very large-scale datasets and an array of diverse parallelization models for fast, high-quality distributed volume rendering of particles, coming from numerical simulations in many file formats, e.g., smoothed particle hydrodynamics from astrophysical simulations.

5.1.3 VD-Maps

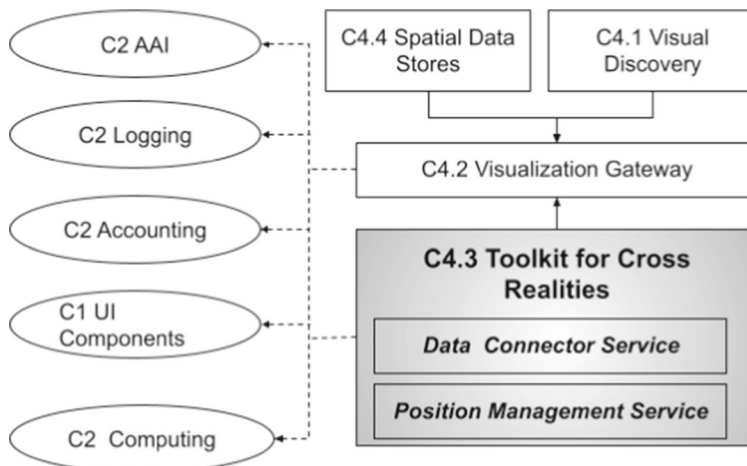
The VD-Maps service consists of a set of tools for visualization of large-scale, high-resolution maps, either 2D (images), 2.5D (elevation/bathymetric maps) or 3D (3D meshes), provided in real-time through a web application. A conversion service is included to translate standard formats into OGC compliant hierarchical data structures through a web visualization viewer based on CesiumJS.⁴⁷

5.2 C4.2 - Visualization Gateway (VG)

The VG service provides a flexible test bed for designing, prototyping, realising and fully validating complex visualization solutions for customised data exploration workflows. VG (see Fig. 3) is founded on

⁴⁷<http://www.cesium.com/cesiumjs/>

Fig. 4 NEANIAS C4.3 Toolkit for Cross Realities services diagram and integration with base core services



the popular Jupyter Hub⁴⁸ project which is chosen for its ability to spawn notebook servers in the Cloud, thus streamlining the usage for end users by avoiding manual installation of software tools locally and eliminating the necessity for complex software dependencies and housekeeping mechanisms. The underlying philosophy of the Visualization Gateway is to focus on usability and seamless embedding within the relevant end-user workflows and to interlink not only with C4.1 services but also with other NEANIAS core services, e.g. C3 for visualization of AI-powered solutions, C4.3 to underpin powerful virtual reality solutions and C4.4 to facilitate end-user data accessibility. The Visualization Gateway is deployed in a way that can be fully embedded within the relevant workflows of end-user community activities, while exploiting a range of diverse parallelisation models and infrastructure accelerator capabilities for optimal performances.

5.3 C4.3-Toolkit for Cross Reality (XR)

C4.3 provides a toolkit underpinning an environment for designing, implementing and validating complex visualization solutions exposed to end-users via Cross Reality (XR) mechanisms, particularly those based on Virtual Reality. To serve as a universal core service for multiple users, popular software frameworks and technologies have been selected based on a set of components used by an existing TLR6 software solution called Astra Data Navigator (ADN), e.g. the frameworks provided by the game engine Unity.⁴⁹ Such

components support interactive data exploration and navigation mechanisms, e.g., for advanced comparisons in multidimensional and multi-frequency datasets for research and public outreach. C4.3 services are built upon extending existing components with a focus on providing services with enhanced realism, precision and usability along a number of aspects relating to enriched user experience covering novel navigation mechanisms and seamless integration of large-scale catalogues.

As shown in Fig. 4, C4.3 exploits the framework for visual discovery developed in C4.1. C4.3 services are envisaged to be interconnected seamlessly with C4.2 to furnish complex visualization workflows and C4.4 to underpin advanced data access. The components implemented so far consist of a Data Connector Service (DCS) for retrieving, in a generic way, data coming from different sources (individual files or databases) and a Positioning Manager Service (PMS) providing the capability to retrieve specific data about the position/rotation of particular objects at specific temporal instants. The latter uses Spice kernels,⁵⁰ and is tailored to astrophysics environments. However, the system will eventually support visualization of 3D data from other NEANIAS community sectors too, such as underwater or atmospheric.

5.4 C4.4 - Spatial Data Stores (DS)

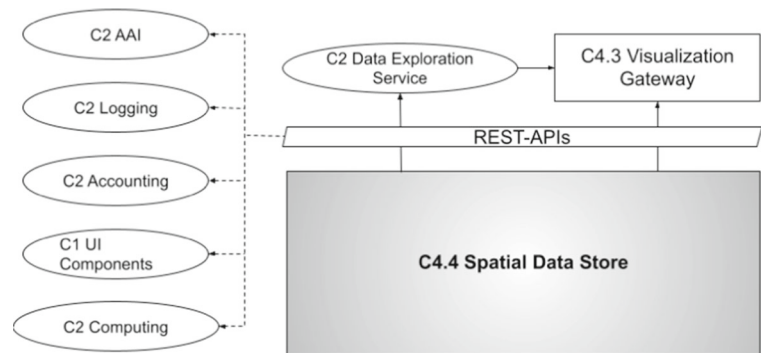
C4.4 services provide a set of reference systems and data structures to facilitate data referencing, access,

⁴⁸<https://jupyter.org/hub>

⁴⁹<https://unity.com/>

⁵⁰<https://naif.jpl.nasa.gov/naif/index.html>

Fig. 5 NEANIAS C4.4 Spatial Data Stores services diagram and integration with base core services



and retrieval by other services, in a standardized manner based on positional/location criteria, e.g. the planet Earth or other planetary bodies (see Fig. 5).

Earth observational data, such as Copernicus datasets and products, are referenced in the ADAM⁵¹ Data Access System (DAS), a software module that manages a large variety of geospatial information supporting multiple data formats and different geographic, geometric and time resolutions. The DAS module allows for accessing, visualizing, subsetting, combining, processing, and downloading multiple data sources at once. DAS also exposes OGC Open Search and Web Coverage Service (WCS 2.x) interfaces that allow for discovering available datasets and subdividing them in any dimension with a single query.

C4.4 services take advantage of Geoserver,⁵² MongoDB,⁵³ PostgreSQL⁵⁴ and PostGIS⁵⁵ systems designed to handle spatial data storage and operations to query/retrieve such data sets from ADAM. ADAM provides seamless full data cycle management functionalities to explore the spatial distribution and temporal evolution of various geophysical and geospatial information, as well as to integrate and execute data processing functionalities at scale. MongoDB, on the other hand, is designed to work with unstructured data, text, numbers and vectors; it has applicability on spatial vector data - data representing polygons are used to e.g. define observation footprints. Geoserver comes as a good alternative for small size, heterogeneous data sets, as a volatile data storage system to publish

high-level or temporary user products. PostGIS provides support for geographic objects allowing location queries to be run in plain SQL, following the OGC's guidelines. PostGIS enjoys further wide support from various third-party open source and proprietary tools, such as QGIS.⁵⁶ The Spatial Data Store is a composition of physical resources (i.e. NAS, S3 buckets⁵⁷) and tools (MongoDB, PostgreSQL/PostGIS, and GeoServer) organized to provide OGC features and raster data through C2 data exploration, provided by the ADAM platform, as well as low-level REST based programming interface mechanisms.

6 Service Delivery and Evaluation

A NEANIAS service represents one or more software components that define a set of active resources i.e. resources that answer user requests [17]. Services follow the single responsibility principle, i.e. being specialized in their role and responsible for a single and typically complex task, often requested through a simple interface by the end-users. The software underpinning a service is abstracted from users to experience its functionality in a high-level interaction setting. In the context of a distributed architecture, as it is typically used with cloud computing, interfaces are often provided by URLs. Documentation provides users with information and training about service capabilities and its usage. In NEANIAS, services are built on top of TRL6 software solutions – i.e. software fully functional in its original domain – intended to evolve towards TRL8, in our context, fully operational in a cloud setting.

⁵¹<https://adamplatform.eu>

⁵²<http://geoserver.org/>

⁵³<https://www.mongodb.com/>

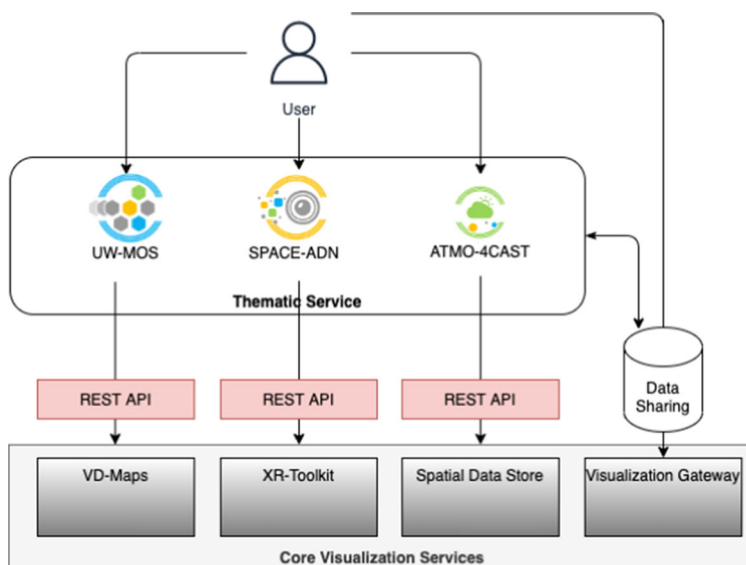
⁵⁴<https://www.postgresql.org>

⁵⁵<https://postgis.net/>

⁵⁶<https://www.qgis.org/en/site/>

⁵⁷<https://aws.amazon.com/s3/>

Fig. 6 NEANIAS service inter-operability scheme, showing the underpinning of thematic services by core services



6.1 Thematic Services Underpinning

As described in Section 4, NEANIAS core services underpin the thematic services that enable specific domain tasks. Therefore, core services are generally not delivered directly to the end users (except for the AI and visualization gateway), but instead, they are consumed by the thematic services for multiple purposes. To ensure the compatibility and interoperability of the different services that compose the NEANIAS ecosystem, inter-service communication is handled through service *contracts*, in which each service defines its interface, resources and interaction patterns. In NEANIAS, service contracts (i) are, whenever possible, REST compliant, (ii) enforce backwards compatibility among versions unless noted otherwise, and (iii) follow widely adopted standards. Such a standardization applies not only to the definition of service operations –e.g., endpoint naming conventions–, but also data models and schemas, to ensure reusability and interoperability.

Figure 6 shows a concrete example of the service interoperability scheme, in which three thematic services make use of some of the visualization core services described in Section 5. From left to right, the *Seafloor Mosaicing from Optical Data* (UW-MOS) service relies on the core service VD-maps to generate 2D/3D tilings from seafloor imagery; the *Astra Data Navigator* (SPACE-ADN) seamlessly exploits the XR-toolkit to provide interactive navigation of stellar

catalogues, such as Gaia;⁵⁸ and the *Atmospheric Forecast* (ATMO-4CAST) service makes extensive use of the Spatial Data Store to generate accurate meteorological forecasts.

6.2 Releases and Deployment Management

The development of NEANIAS services follows the Continuous Integration / Continuous Delivery (CI/CD) philosophy, in order to deliver updates and new features to end users in a more frequent and efficient way. Figure 7 depicts the typical lifecycle of a NEANIAS service, from the development of new features to the final release of a new version. This process can be divided into three major phases:

- *development*: in this phase, code changes are implemented, including new features –that extend the functionality of the service–, refactorings and bug fixes. New features are planned regularly, according to the original design requirements of the service, and the user feedback received.
- *validation*: in this phase, the service is exhaustively evaluated to ensure that it works as expected and the design requirements are met. Validation comprises two aspects: testing and end-user assessment. Testing can be performed automatically to reduce bug finding costs. The employed CI/CD

⁵⁸<https://gea.esac.esa.int/archive/>




Fig. 7 NEANIAS Services Release process steps

pipelines include unit and integration testing stages, as provided by the NEANIAS Gitlab. On the other hand, end-user assessment is performed prior to release by groups of expert domain users, and involves: a) verifying of the correct implementation, functionality and performance of each service component; b) verifying that communications among components are performed correctly; and, finally, c) end-to-end testing to simulate end-user and administrator behaviours on the platform. Finally, at this stage, comprehensive documentation is produced in order to facilitate the usage of the service.

- *publishing*: this phase starts with the effective deployment of the service. For deployment, NEANIAS services follow the development, staging and production model. First, newly implemented code and features are deployed to a development server, where basic testing takes place to spot bugs, application misbehaviours and other potential issues. The most intensive testing occurs in the staging server, by running comprehensive test suites to detect any remaining bugs, gathering performance metrics and finding whatever needs to be polished prior to final deployment to the production server, where the service is made available for real-world end users. Before this final stage, it is customary for the service provider to run a *smoke test*, covering the minimal set of functionalities required for the application to work at a minimum, thus assessing its stability and allowing for a quick rollback if necessary. After a first production deployment, NEANIAS services are published on the NEANIAS service catalogue,⁵⁹ and consequently may be onboarded in the EOSC marketplace.⁶⁰

A released service is defined chiefly by two aspects: its documentation and interface, and its access

endpoint, i.e. the point through which it can be accessed. The way a service is provided – in other words its interface to prospective users – varies; some services collect information from others (e.g., for log aggregation), others answer to queries (e.g. spatial data stores), while others are less dynamic, openly providing data content or guidelines (e.g. the web toolkit). Regardless of the nature of a particular service or its interface, its documentation is nevertheless fundamental, even for base core services. The idea is at all times to state clearly and persistently how to properly use the service, and list comprehensively all necessary dependencies required for its operation.

6.3 Service Evaluation

We are following a robust technical evaluation protocol for all thematic services (and their underlying core services) in NEANIAS, and in particular those providing visualization capabilities. This involves an evaluation cycle for each service release (see Section 7) manifested in real-world production environments, and carried out by cohorts of expert community users from the different science domains: space, see [16]; underwater, see [26]; and atmospheric, see [23].

Starting from thematic activities NEANIAS is now embarking on a larger scale service validation by external users from other community bodies for developing further improvements and new functionalities in preparing the ground for EOSC onboarding on top of the core visualisation services while always ensuring their generic flavour.

7 Service Release Status

The first release made services available under the NEANIAS development infrastructure [17]. Currently all services are in their second release iteration with a final delivery expected by the end of NEANIAS - see below for further details.

⁵⁹<https://catalogue.neanias.eu/>

⁶⁰<https://marketplace.eosc-portal.eu/>

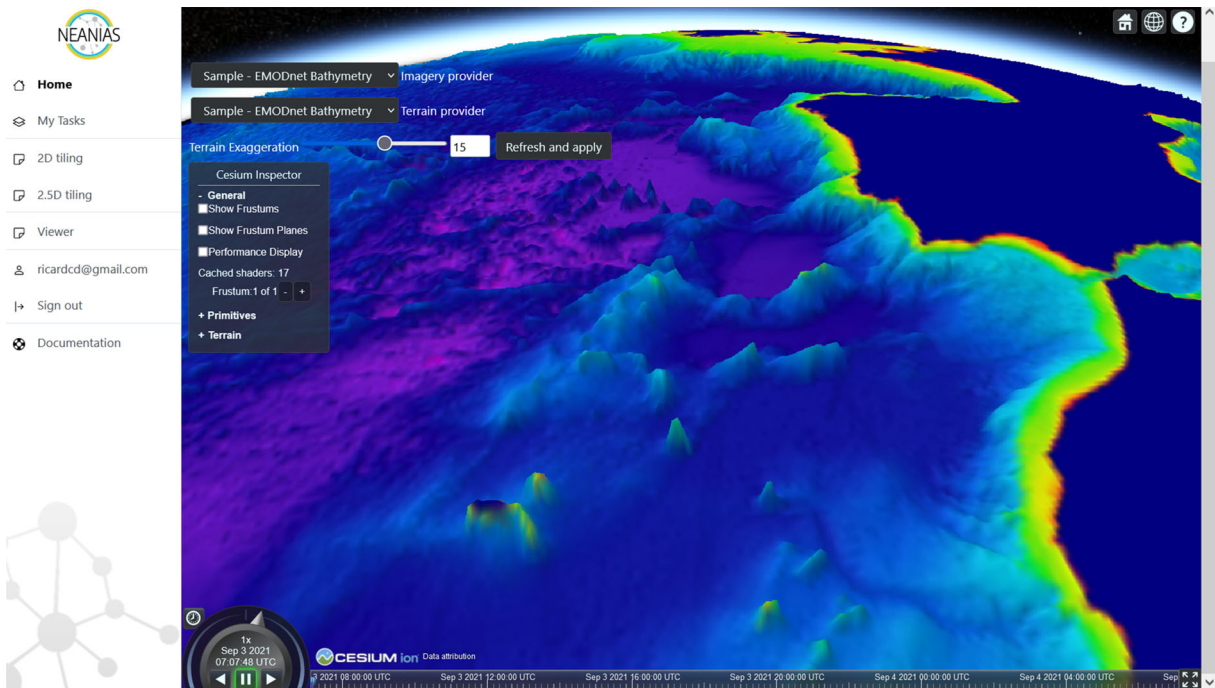


Fig. 8 NEANIAS VD-Maps Service to process and visualize imagery for terrain-based hierarchical tiled maps

VD-VisIVO is released as a Docker container⁶¹ managed through GitLab⁶² and delivered via the NEANIAS container image registry. The source code of VisIVO is freely available on GitHub,⁶³ licensed under GNU, and includes demos and user manuals. Similarly, VD-Splotch⁶⁴ is also released as a Docker container with software implementation managed through GitLab⁶⁵ and delivery to end-users via the NEANIAS container image registry.

Finally VD-Maps (see Fig. 8) is released as a Docker container providing additional containers for serving and visualizing the generated maps. The code and instructions to build these containers are available through NEANIAS GitLab,⁶⁶ the docker images are available in the registry while external dependencies such as the Cesium Terrain Builder⁶⁷ and EMODnet

Quantized Mesh Generator⁶⁸ libraries are hosted separately in their respective repositories.

The visualization gateway⁶⁹ is based on Jupyter Hub⁷⁰ - its novelty is on showcasing a FAIR platform integrating distinct visualization approaches into an ecosystem allowing realization of complex scenarios through seamless integration mechanisms with common scientific workflows. The gateway was prototyped using KinD,⁷¹ originally in a local setting and prior to releasing within the GARR Container Platform.⁷² The web interfaces are accessible via NEANIAS AAI and offer separate notebook spawners for each of the Visual Discovery Framework components - these can be accessed on the notebook through interactive python demonstrations (see Fig. 9). Prior to spawning a notebook, the system offers the option

⁶¹<https://www.docker.com>

⁶²<https://gitlab.neanias.eu/c4-service/vd-visivo>

⁶³<https://github.com/inaf-oact-VisIVO/VisIVOServer>

⁶⁴<https://github.com/splotchviz/splotch>

⁶⁵<https://gitlab.neanias.eu/c4-service/vd-splotch>

⁶⁶<https://gitlab.neanias.eu/c4-service/vd-maps>

⁶⁷<https://github.com/geo-data/cesium-terrain-builder>

⁶⁸<https://github.com/coronis-computing/emodnet-qmgc>

⁶⁹<https://vis-gateway.neanias.eu/>

⁷⁰<https://jupyter.org/hub>

⁷¹<https://kind.sigs.k8s.io>

⁷²<https://cloud.garr.it>

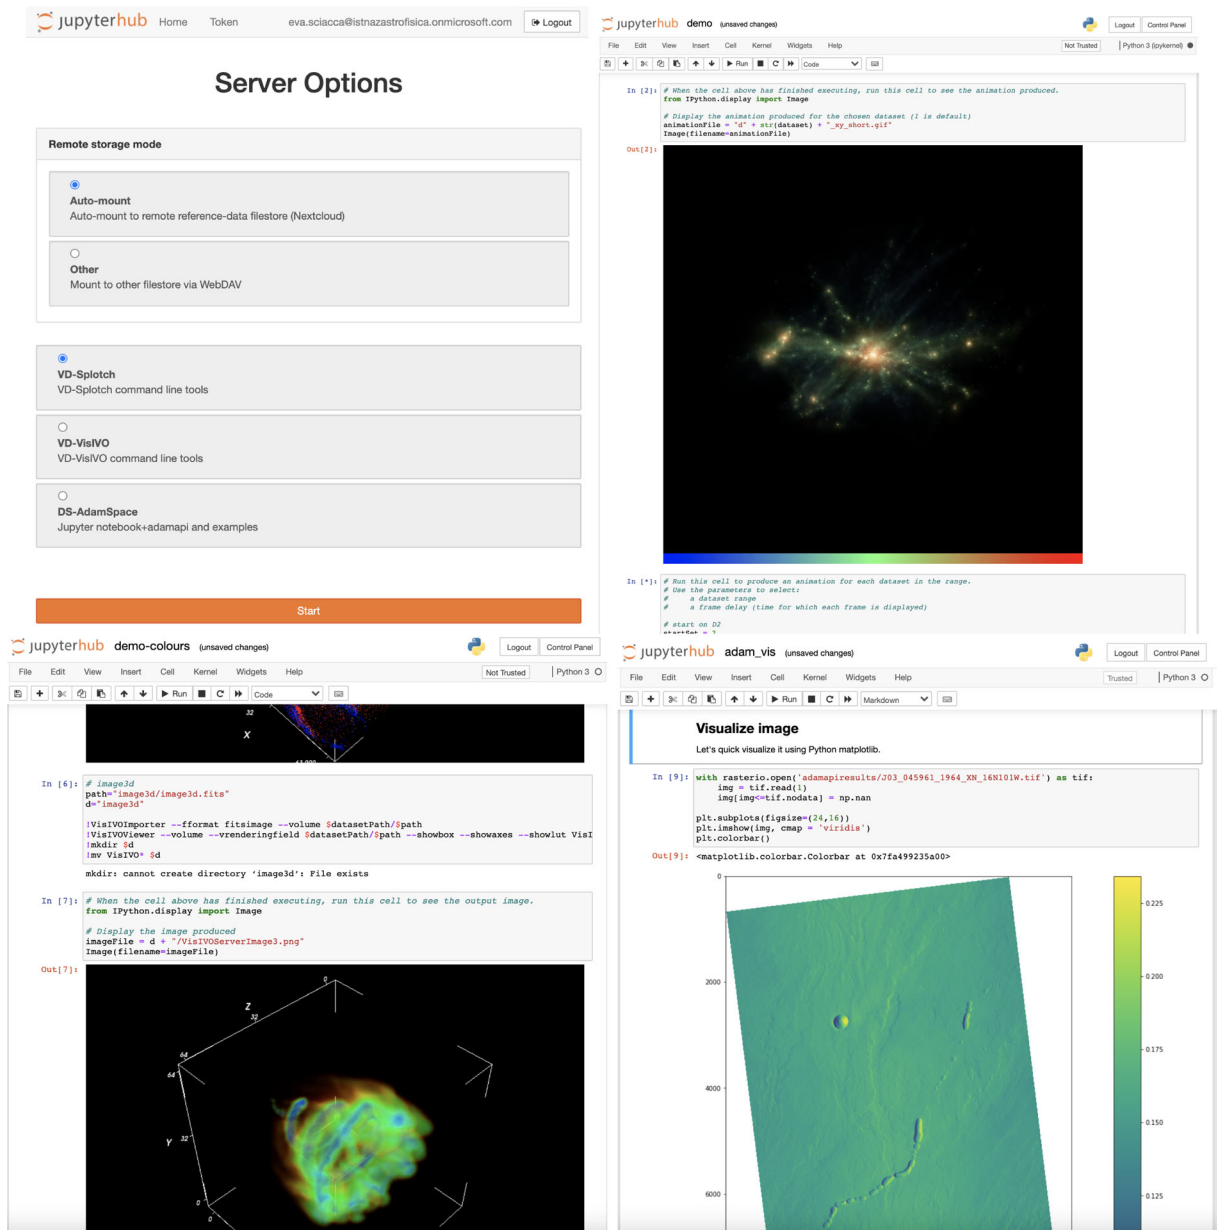


Fig. 9 The NEANIAS Visualization Gateway main page (top left) showing sample notebooks integrating VD-Splotch (bottom left) and Spatial DS (bottom right)

to mount the notebook to a Nextcloud⁷³ store via WebDAV.⁷⁴ For the purposes of demonstrating the basic functionality of the tools, a notebook mounts by default to the NEANIAS Nextcloud store which contains read-only reference datasets. Alternatively,

⁷³<https://nextcloud.com>

⁷⁴<http://www.webdav.org>

users can specify credentials for a different Nextcloud store, or make use of the file uploading capability of the notebook web interface to manually upload their own datasets. The authentication service based on Keycloak⁷⁵ allows notebook access to authorized users via the Jupyter Hub login page. Integration with

⁷⁵<https://www.keycloak.org>

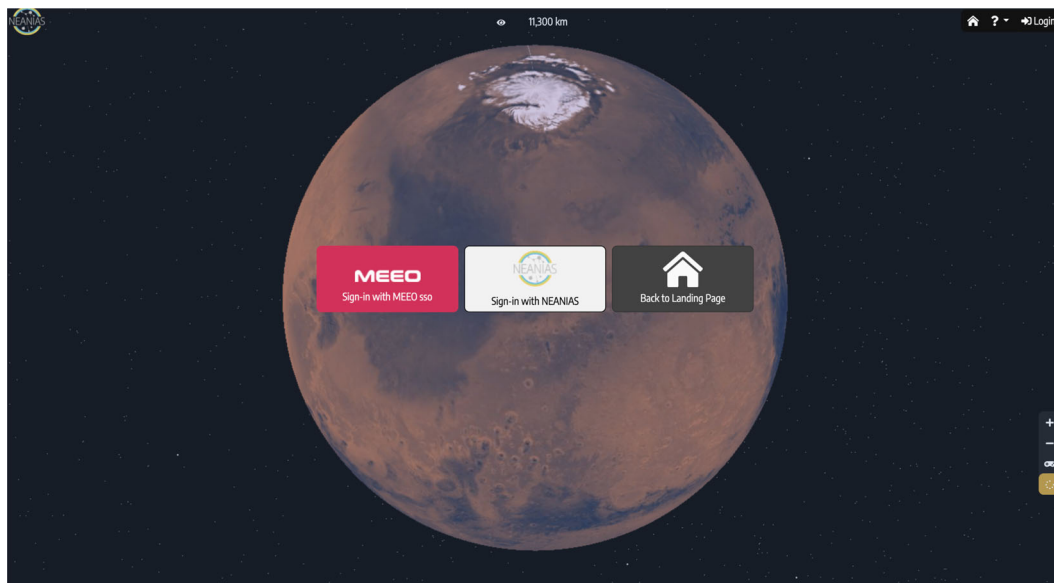


Fig. 10 NEANIAS ADAM-Space platform exploring data from Mars underpinned by the Spatial Data Store service C4.4

NEANIAS accounting and logging services, using filebeat⁷⁶ to pass on Kubernetes logs, is currently underway. This will allow increasing reliability of the system as well as accounting for global usage of resources thus optimizing the current setup. Also linkages with Zenodo, toward Open Science integration [12, 22], is being tested for ensuring results reproducibility.

The XR toolkit has released the PMS⁷⁷ and DCS⁷⁸ modules enabling ADN to extend navigation mechanisms within the Hipparcos catalogue [13] (~110K celestial objects) with the ultimate goal of providing an interactive exploration solution for catalogues such as Gaia[5] (~1.8 billion objects) and beyond. The novelty of our solution is in stark contrast to existing VR solutions, (e.g. CosmoScout VR⁷⁹ or SpaceEngine⁸⁰) which allow representation of stellar catalogues of limited sizes only, and are developed entirely from scratch (i.e. not based on robust game engine platforms), thus making cumbersome

their long term maintenance and sustainability. The PMS module offers a simplified utilization of the Spice toolkit⁸¹ to retrieve the correct positions for all celestial objects present in rendered catalogue scenes - this is achieved by exploiting appropriate request parallelization mechanisms for optimizing overall performances.

The spatial data store service is provided by the MEEO Cloud storage resources⁸² to host atmospheric and planetary science data. An OpenNebula⁸³ instance manages the MEEO Data Center virtualization resources. Kubernetes is used for serving discovery and access requests. The datasets and products metadata are archived in a catalogue, and original products are archived in physical volumes. The DS service is accessed through the space service ADAM-Space⁸⁴ for exploration of data related to Mars (see Fig. 10) and from the ADAM-Platform⁸⁵ for exploration of environmental data of the Earth.

⁷⁶<https://www.elastic.co/guide/en/beats/filebeat/current/filebeat-overview.html>

⁷⁷<https://pms.neanias.eu/>

⁷⁸<https://dcs.neanias.eu/>

⁷⁹<https://github.com/cosmoscout/cosmoscout-vr>

⁸⁰<http://spaceengine.org/>

⁸¹<https://naif.jpl.nasa.gov/naif/toolkit.html>

⁸²www.meeo.it/cloud

⁸³<https://opennebula.io>

⁸⁴<https://explorer-space.adamplatform.eu/>

⁸⁵<https://explorer.adamplatform.eu/>

8 Conclusions

The NEANIAS project addresses the prototyping of new innovative solutions in EOSC, driving the co-design, delivery, and integration of thematic and core services, derived from state-of-the-art research assets and practices in three major sectors: underwater, atmospheric and space research. The underlying ambition is to create a collaborative research ecosystem promoting Open Science principles and practices, injecting add-on value to the data and contributing to the effective materialization of the EOSC landscape services.

As a cornerstone of scientific discovery, data visualization is crucial for unlocking new discoveries and generating knowledge. The big data paradigm sweeping in almost every scientific discipline requires innovative and efficient data visualization solutions. We have delved in this paper into the details of NEANIAS core visualization services. We have summarized the scientific challenges to overcome in the different research communities, the role of core services within the NEANIAS system architecture, and the main features provided by each visualization service, specifically on visual discovery, visualization gateway, cross-reality and spatial data stores. Finally, we have reviewed the service delivery model and the current software release status of the visualization core services.

The development of NEANIAS services is on going and is envisaged to be completed by the end of 2022. CI compliant progressive releases are being planned, taking into account existing and newly gathered user requirements. These releases will eventually push the NEANIAS visualization services to TRL8, promoting open science practices by fully onboarding the entire service portfolio into the marketplace of the EOSC ecosystem.

Funding Open access funding provided by INAF - National Institute for Astrophysics within the CRUI-CARE Agreement.

Data Availability Statement Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

Compliance with Ethical Standards

Conflict of Interests The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ahrens, J., Geveci, B., Law, C.: Paraview: An end-user tool for large data visualization. *The Visualization Handbook* 717(8) (2005)
- Ayachit, U., Bauer, A., Geveci, B., O'Leary, P., Moreland, K., Fabian, N., Mauldin, J.: Paraview catalyst: Enabling in situ data analysis and visualization. In: *Proceedings of the First Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization*, pp. 25–29 (2015)
- Becciani, U., Sciacca, E., Costa, A., Massimino, P., Pistagna, C., Riggi, S., Vitello, F., Petta, C., Bandieramonte, M., Krokos, M.: Science gateway technologies for the astrophysics community. *Concurr. Comput.: Pract. Exp.* **27**(2), 306–327 (2015)
- Bordiu, C., Bufano, F., Sciacca, E., Riggi, S., Molinaro, M., Vizzari, G., Krokos, M., Brandt, C.: Astronomical research in the next decade: Trends, barriers and needs in data access, management, visualization and analysis. *Astronomical Data Analysis Software and Systems XXX* proceedings, to be published (2020)
- Brown, A., Vallenari, A., Prusti, T., De Bruijne, J., Babusiiaux, C., Bailer-Jones, C., Biermann, M., Evans, D.W., Eyer, L., Jansen, F., et al: Gaia data release 2-summary of the contents and survey properties. *Astron. Astrophys.* **616**, A1 (2018)
- Campos, R., Quintana, J., Garcia, R., Schmitt, T., Spoelstra, G., Schaap, M.A.D.: 3d simplification methods and large scale terrain tiling. *Remote Sens.* **12**(3), 437 (2020)
- Dias, D., Pina, N., Tchepel, O.: Characterization of traffic-related particulate matter at urban scale. *Int. J. Transp. Develop. Integr.* **3**(2), 144–151 (2019)
- Feng, X., Shen, J., Fan, Y.: Rest: An alternative to rpc for web services architecture. In: *2009 First International Conference on Future Information Networks*, pp. 7–10. IEEE (2009)
- Franke, L., Haehn, D.: Modern scientific visualizations on the web. *Informatics* 7(4) (2020). <https://doi.org/10.3390/informatics7040037>. <https://www.mdpi.com/2227-9709/7/4/37>
- Goodale, T., Allen, G., Lanfermann, G., Massó, J., Radke, T., Seidel, E., Shalf, J.: The cactus framework and toolkit: Design and applications. In: *International Conference on*

- High Performance Computing for Computational Science, pp. 197–227. Springer (2002)
11. Kacsuk, P., Farkas, Z., Kozlovsky, M., Hermann, G., Balasko, A., Karoczkai, K., Marton, I.: Ws-pgrade/guse generic dci gateway framework for a large variety of user communities. *J. Grid Comput.* **10**(4), 601–630 (2012)
 12. Mendez, K.M., Pritchard, L., Reinke, S.N., Broadhurst, D.I.: Toward collaborative open data science in metabolomics using jupyter notebooks and cloud computing. *Metabolomics* **15**(10), 1–16 (2019)
 13. Perryman, M., Lindegren, L., Kovalevsky, J., Hog, E., Bastian, U., Bernacca, P., Creze, M., Donati, F., Grenon, M., Grewing, M., et al: The hipparcos catalogue. *A&A* **500**, 501–504 (1997)
 14. Raji, M., Hota, A., Hobson, T., Huang, J.: Scientific visualization as a microservice. *IEEE Trans. Visual. Comput. Graph.* **26**(4), 1760–1774 (2018)
 15. Rivi, M., Gheller, C., Dykes, T., Krokos, M., Dolag, K.: Gpu accelerated particle visualization with splotch. *Astron. Comput.* **5**, 9–18 (2014)
 16. Rossi, A.P., Brandt, C.H., et al.: NEANIAS Deliverable D4.4 Report on the Developed and Validated Space Thematic Services #1. Tech. rep. H2020 NEANIAS Project (2020)
 17. Rossi, A.P., Brandt, C.H., et al.: NEANIAS Deliverable D6.3 Core Services Software Release Report. Tech. rep. H2020 NEANIAS Project (2020)
 18. Sciacca, E., Becciani, U., Costa, A., Vitello, F., Massimino, P., Bandieramonte, M., Krokos, M., Riggi, S., Pistagna, C., Taffoni, G.: An integrated visualization environment for the virtual observatory: Current status and future directions. *Astron. Comput.* **11**, 146–154 (2015)
 19. Sciacca, E., Kakalettris, G., et al.: NEANIAS Deliverable D6.1 Core Services Architecture, Design Principles and Specifications. Tech. rep. H2020 NEANIAS Project (2020)
 20. Sciacca, E., Krokos, M., Ugo, B., Bordiu, C., Bufano, F., Costa, A., Pino, C., Riggi, S., Vitello, F., Brandt, C., et al: Novel EOSC services for Space Challenges: The NEANIAS First Outcomes. In: *Astronomical Data Analysis Software and Systems XXX* proceedings, to be published (2020)
 21. Sciacca, E. et al.: NEANIAS Deliverable D6.4 Core Services Architecture, Design Principles and Specifications (update). Tech. rep. H2020 NEANIAS Project (2021)
 22. Sicilia, M.A., García-Barriocanal, E., Sánchez-Alonso, S.: Community curation in open dataset repositories: Insights from zenodo. *Procedia Comput. Sci.* **106**, 54–60 (2017)
 23. Tibaldi, A. et al.: NEANIAS Deliverable D3.4 Report on the Developed and Validated Atmospheric Thematic Services #1. Tech. rep. H2020 NEANIAS Project (2020)
 24. Wells, D.C., Greisen, E.W., Harten, R.H.: FITS - a Flexible Image Transport System **44**, 363 (1981)
 25. Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.W., da Silva Santos, L.B., Bourne, P.E., et al: The fair guiding principles for scientific data management and stewardship. *Scientific Data* **3**(1), 1–9 (2016)
 26. Wintersteller, P. et al.: NEANIAS Deliverable D2.4 Report on the Developed and Validated Underwater Thematic Services #1. Tech. rep. H2020 NEANIAS Project (2020)
 27. Zhou, Y., Weiss, R.M., McArthur, E., Sanchez, D., Yao, X., Yuen, D., Knox, M.R., Czech, W.W.: Webviz: A web-based collaborative interactive visualization system for large-scale data sets. In: *GPU Solutions to Multi-scale Problems in Science and Engineering*, pp. 587–606. Springer (2013)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.