

D4Science
Data Infrastructure
Service Level Agreement

Service Level
Agreement Data

Customer
Provider
Start Date
End Date
Status
Agreement Date

OpenMod Initiative
D4Science.org
TBC
TBC
DRAFT
TBC

CHANGE LOG

Reason for change	Issue	Actor	Date
Proposed Service Level Agreement	1.0	D4Science.org	30/08/2019

DRAFT

LIST OF ABBREVIATIONS

ABAC	Attribute-based access control
D4Science	Distributed infrastructure for collaborating communities
IT	Information Technology
TLS	Transport Level Security
VRE	Virtual Research Environment
SDI	Spatial Data Infrastructure

DRAFT

TABLE OF CONTENTS

CHANGE LOG	2
LIST OF ABBREVIATIONS.....	3
TABLE OF CONTENTS	4
LIST OF TABLES	6
1 SLA COORDINATES.....	7
2 D4SCIENCE INFRASTRUCTURE SECURITY.....	9
3 THE SERVICES	10
4 THE D4SCIENCE SERVICES TERMS OF USE	12
4.1 Virtual Research Environment Management Service.....	12
4.2 Workspace Service.....	12
4.3 Social Networking Collaborative Platform	13
4.3.1 Messaging Service.....	14
4.3.2 Notification Service.....	15
4.4 Catalogue Service	15
4.5 Data Analytics Platform.....	17
4.6 Tabular Data Manager (TDM) Platform.....	19
4.7 Spatial Data Infrastructure (SDI)	21
4.8 Gateway Service.....	23
5 VRE SPECIFIC SERVICES TERMS OF USE.....	25
5.1 Share Latex (fully integrated)	25
6 SERVICE CAPACITIES	26
7 SERVICES HOURS AND EXCEPTIONS	28
8 SERVICE LEVEL TARGETS.....	29
9 SERVICE LEVEL COMPENSATION.....	30
9.1 Training Module Cost.....	32
10 SUPPORT	34
10.1 Incident handling	34
10.2 Service requests	34
10.3 Service Operation	34
11 LIMITATIONS AND CONSTRAINTS	36
12 COMMUNICATION, REPORTING AND ESCALATION	37
12.1 General communication.....	37
12.2 Regular reporting.....	37
12.3 Escalation and complaints.....	37

13 INFORMATION SECURITY AND DATA PROTECTION.....38

14 RESPONSIBILITIES.....39

 14.1 Of the Provider39

 14.2 Of the Customer.....40

15 REVIEW, EXTENSIONS AND TERMINATION.....41

16 GENERAL PROVISIONS.....42

 16.1 Non-assignability42

 16.2 Entire Agreement.....42

 16.3 Governing Law.....42

 16.4 Force Majeure.....42

DRAFT

LIST OF TABLES

Table 1. Service Level Targets	29
Table 2. Compensation Costs per Bundle	31
Table 3. Maintenance and Operation Effort in Working Days per Bundle	32
Table 4. Response times to incidents according to the incident priority of "Medium" services table	34
Table 5. Service Operation Procedures and Rules	34
Table 6. Contacts	37
Table 7. Reporting	37

DRAFT

1 SLA COORDINATES

The present Service Level Agreement ("the Agreement") is made between the
D4Science.org, legally represented by CNR (the Provider)
and
Openmod Initiative (the Customer),

legally represented by _____

to define the **provision** of the OPENMOD VRE equipped with the services described hereafter.

Representatives and contact information are defined in Section 12.

With the term provision it is intended the deployment and operation of the **dedicated** Virtual Research Environment.

Summary coordinates:

VRE Name: OPEMOD

VRE Description:

The OPENMOD VRE _____

VRE Access policies: **Restricted**

One of the following VRE access policies can be selected at any time during the operation of the VRE:

- **Private access** policy: the VRE is private and a user can access it only by invitation issued by the VRE Manager. Upon acceptance of the user of the invitation, the user becomes member of the VRE with a user role;
- **Restricted access** policy: a user registered to the gateway and with a valid identity can request access to the VRE. The VRE Manager can approve or reject any user request. In case of approval, the user becomes member of the VRE with a user role;
- **Public access** policy: a user registered to the gateway and with a valid identity can become member of the VRE by simply accessing it. The VRE Manager is just notified.

Four different roles are supported by default and additional ones can be defined as per request of the Customer:

- **VRE Manager** role: any user with this role can manage user registration/deregistration to the VRE and assign/remove roles by accessing the Administration dashboard.
- **User** role: any user with this role can access the VRE data and share private data with other selected members of the VRE;
- **Processor** role (Optional): any user with this role can register its own process and execute a process by accessing the Method Engine service;
- **Editor** role: any user with this role can publish data to all members of the VRE by exploiting the Catalogue service.

This Agreement is valid from _____ to _____ and it can be renewed up to 5 years from the signature at the condition reported in this document. After 5 years from the signature of this SLA, the cost per month for the services can be varied and increased / decreased to an extent not exceeding 15% of what is currently established.

The Agreement was discussed and approved by the Customer and the Provider on _____
(DATE).

2 D4SCIENCE INFRASTRUCTURE SECURITY

D4Science.org provides access to a set of services, as reported in the following sections, hosted by three different sites in the EU.

The connection between the sites is **secured with Transport Level Security** (TLS) that provides communications security over the computer network.

D4Science.org ensures privacy and data integrity between two communicating computer applications. In particular, any connections between a client (e.g., a web browser) and a D4Science.org server have the following properties:

- The connection is *private* (or *secure*) thanks to the adoption of the symmetric cryptography to encrypt the data transmitted. The keys for this symmetric encryption are generated uniquely for each connection and are based on a shared secret negotiated at the start of the session. The server and client negotiate the details of which encryption algorithm and cryptographic keys to use before the first byte of data is transmitted. The negotiation of a shared secret is both secure (the negotiated secret is unavailable to eavesdroppers and cannot be obtained, even by an attacker who places themselves in the middle of the connection) and reliable (no attacker can modify the communications during the negotiation without being detected);
- The identity of the communicating parties can be *authenticated* using public-key cryptography. This authentication can be made optional at client side, but is ensured at the server side;
- The connection ensures *integrity* because each message transmitted includes a message integrity check using a message authentication code to prevent undetected loss or alteration of the data during transmission;
- The connection ensures forward secrecy, ensuring that any future disclosure of encryption keys cannot be used to decrypt any TLS communications recorded in the past.

D4Science.org provides access to services and data via **Virtual Research Environments** (VREs). Each VRE enables services and data exploitation to the users authorized to access the VRE.

D4Science.org authorization is empowered by a **token-based authorization system** compliant with the Attribute-based access control (ABAC) that defines an access control paradigm whereby access rights are granted to users through the use of policies that are validated in a VRE context.

3 THE SERVICES

D4Science provides researchers and practitioners with a working environment where Open Science practices are supported and promoted: **better interpretation, understanding and reproducibility of research activities and results** and **reduction of the overall cost of research by promoting re-use of results** are enabled through the provision of a truly integrated set of services.

D4Science is designed to deliver tailored working environments to support the needs of its designated communities through the provision of Virtual Research Environments (VREs). **Each VRE is equipped with at least the basic services supporting collaboration and cooperation among its users**, namely: (i) a *shared workspace* to store and organise any version of a research artefact; (ii) a *social networking area (enabled by three services: messaging, social networking, and notification)* to have discussions on any topic (including working version and released artefacts) and be informed on happenings; (iii) a *data analytics platform* to execute processing tasks either natively provided by VRE users or borrowed from other VREs to be applied to VRE users' cases and datasets; and (iv) a *catalogue-based publishing platform* to make the existence of a certain artefact public and disseminated. To this set of services, others can be added to manage specific data types as tabular datasets and spatial datasets.

Moreover, D4Science is a trusted service provider of the European Open Science Cloud (EOSC). EOSC aims to give the Union the full benefits of data-driven science, by offering "1.7 million European researchers and 70 million professionals in science and technology a virtual environment with free at the point of use, open and seamless services for storage, management, analysis and re-use of research data, across borders and scientific disciplines".

The OPENMOD Virtual Research Environment is equipped with the following facilities:

- A **VRE Management service** to enable authorized users (i.e. VRE Managers) to manage other users using or wanting to access the VRE. VRE Managers can (i) authorize users for access to the VRE, (ii) assign or withdraw roles to users, (iii) remove users, and (iv) send communications to the current users;
- A shared **workspace service** to enable every user to store and organise the information objects he/she is interested in working with. In addition to that, the user is allowed to collaborate with other users by sharing objects and messages. To every user is guaranteed 100 GB storage volume;
- A **social networking collaboration platform** to enable users to use the common facilities typical of social networks – e.g., posting news, commenting on posted news – yet adapted to the settings of the working environments. Users can post news as well as applications. This platform embeds the following services:
 - A **messaging service** to provide users with a common email environment as-a-Service. The distinguishing feature is represented by its integration with the other services, e.g., it is possible to send any information object residing in the workspace (regardless of how "big" and "complex" it may be) as an attachment without consuming bandwidth;
 - A **notification service** to alert users on relevant activities as they happen. These notifications offer a sense of anticipation and create a productivity boost. Users receive an alert (through a priori selected channels, e.g., email, web portal, twitter) notifying them when something of interest has happened in their VRE(s); The notification service is equipped with a **members facility** to provide users with a list of the VRE co-workers, i.e. the list of members partaking in the VRE and contributing to it;
- A **catalogue service** to provide users with the capability to search, browse, discover, and filter any resource, e.g. datasets and methods, published in the catalogue by VRE members with the editor role;
- A **gateway service** to provide users with a web portal to access the VREs.

To this mandatory set of facilities, the following ones may be added:

- A **data analytics platform** to provide users with the capability either to integrate their own methods or to exploit already integrated methods in the VRE. The methods can be implemented in Java, Python, R, Octave, Matlab. The version of the platform, e.g. the Python version, must be the same in a VRE but may be different in different VREs;
- A **tabular data manager platform** to provide users with a collaborative toolset to collate, harmonize and manage tabular data (e.g. fishery time series, code lists); to produce datasets relevant to the ecosystem approach to fisheries and the conservation of marine living resources. The harmonization includes conversion from private formats to international formats such as SDMX and access to environment where it is possible to use R-programming;
- A **spatial data infrastructure** set of services to proper manage geo-referenced data and metadata;
- A **ShareLatex** instance to deliver an online LaTeX editor that allows real-time collaboration and online compiling of projects to PDF format.

DRAFT

4 THE D4SCIENCE SERVICES TERMS OF USE

4.1 Virtual Research Environment Management Service

In D4Science.org, services and related access to data are available through Virtual Research Environments (VREs). These VREs can take various forms including web interactive user interface, web applications, and pluggable standalone user interface.

VREs can be created for a specific period and a specific task and only authorized users access data and services exposed through these VREs.

There are different levels of user authorization available:

- VRE Managers have the right to grant access to other users and to assign roles;
- VRE Editors, i.e. VRE user with Editor role, have the right to register/unregister datasets and make them available through the VRE Catalogue;
- VRE Processors (Optional), i.e. VRE users with Processor role, have the right to register/unregister software methods, and make them available to all Users of the VRE;
- VRE Users can upload, share and manage data within their VRE(s) according to their VRE user profile. They can request access to any Restricted VRE.

Users can apply to a **Private, Restricted** or a **Public** Virtual Research Environment:

- Access to the Private VREs requires a valid user registration performed through the Gateway service. The acceptance of an invitation issued by a VRE Manager is solely under the responsibility of the VRE User that has to verify and accept the VRE Terms of Use;
- Access to the Restricted VREs requires a valid user registration performed through the Gateway service. The acceptance in a VRE of an application issued by a user is solely under the responsibility of the VRE Manager that has to verify that the user belongs to the community operating the VRE;
- Access to the Public VRE is regulated by an automatic procedure that grants the access to any registered user applied for it.

By exploiting a VRE, any user shall be deemed to accept the conditions of use reported in the following sections.

4.2 Workspace Service

The Shared Workspace is an online environment to support secure and controlled data storage and sharing.

This facility relies on Apache Jackrabbit for storing and managing workspace item representations. Items payload are stored by relying on a hybrid cloud storage solution that, by means of ad-hoc plugins, exploits various storage solutions suitable for diverse typologies of content, e.g. MongoDB for binary files, GeoServer and THREDDS Data Servers for geospatial data, RDB for tabular data.

Every workspace item is

- equipped with an actionable unique identifier that can be used for citation and access purposes;
- versioned and a new version is automatically produced whenever the item is explicitly changed by a user or any application/service of the VRE on behalf of an authorised user;
- equipped with rich and extensible business metadata that capture descriptive features as well as lineage features
- organised in folders that can be
 - private: content is available only for its owner;
 - shared: content is available for selected users (decided by the owner);
 - VRE folder, content is available to VRE members;

The Shared Workspace is tightly integrated with both the social networking and the catalogue for easing the dissemination of its artefacts.

Data Managers can exploit the Shared Workspace while keeping accounting and traceability. Individual scientists easily share data, experiments, derived data products with selected colleagues. Working groups empower the working group with an easy-to-use environment where operations are traced and accounted.

By accessing the Workspace service the user shall be deemed to accept these conditions of use:

- The user may use the Workspace to store, retrieve, query, serve, and execute content that is owned, licensed or lawfully obtained by you;
- The user can post content to the Workspace only if (a) the user created and own the rights to the content or the user has the owner's express permission to post the content; and (b) the content does not infringe any other person's or entity's rights (including the copyrights, trademarks, or privacy rights) or violate any applicable laws, the D4Science Terms of Use, Privacy Policy, or any other posted policies;
- The ownership of any intellectual property rights is not in any way transferred to the D4Science.org Infrastructure;
- In the case the user decides to share it with other Users, the user remains responsible for any misuse of the data by other Users;
- At any time the user can decide to unregister her data. All requests are immediately accepted and operated by the D4Science.org Infrastructure;
- The D4Science.org Infrastructure can reproduce, modify, and generate derivative works (such as those resulting from translations, adaptations or other changes required by the management capabilities of the D4Science.org Infrastructure);
- The D4Science.org Infrastructure will not otherwise move or distribute your data for any purpose, except when required to do so by law;
- The D4Science.org Infrastructure is not responsible of the data uploaded and hosted by the infrastructure;
- The D4Science.org Infrastructure will not be responsible of any issue regulating intellectual property rights infringement or illegal use of user's data;
- The D4Science.org Infrastructure will make reasonable efforts to ensure that data are persisted. In the event of hardware or software Failures caused by failures to a hard drive or power supply, the D4Science.org Infrastructure will make reasonable attempts to restore your data. No guarantee whatsoever is provided on the success of any user's data recovery;
- If the D4Science.org Infrastructure reasonably believes any of the Content violates the law, infringes or misappropriates the rights of any third party, you will be notified and may request that such content be removed from the Workspace or access to it be disabled. If you do not remove or disable access to the Prohibited Content within ONE business day of our notice, the D4Science.org infrastructure manager may remove or disable access to the Prohibited Content or suspend its access;
- The user can use the data other users share with her/him in any circumstance and for all usage, reproduce the data, modify the data, and make derivative data based upon the original data, communicate to the public, including the right to reproduce or display the data or copies thereof to the public and perform publicly, as the case may be, the data;
- The user cannot in any circumstance and for any usage redistribute the data other users share with her/him or copies thereof, lend and rent those data or copies thereof, sub-license rights in the data or copies thereof.

4.3 Social Networking Collaborative Platform

The Social Networking Collaborative Platform enables a comprehensive and collaborative environment to support sharing and collaboration among users. It resembles a social

networking environment with posts, tags, mentions, comments and reactions, yet its integration with the other services makes it a powerful and flexible communication channel for researchers.

This facility relies on the Social Networking Engine, a Cassandra database for storing social networking related data and on Elasticsearch for the retrieval of social networking data. The Engine exposes its facilities by an HTTP REST Interface and comprises two services: (i) the Social Networking Service that efficiently stores and accesses social networking data (Posts, Comments, Notifications, etc.) in the underlying Cassandra Cluster, and (ii) the Social Networking Indexer Service that builds Elasticsearch indices to perform search operations over the social networking data.

There is no predefined way to structure a discussion; users can start new discussion threads, annotate them with tags for easing the cataloguing and discovery, refer to other threads and material both internally stored and available on the web. Every (re)action performed by a user – be it a new post, a reply to a post, or the rating of a certain post or post reply – is carefully captured, documented, and equipped with an actionable unique identifier that can be used for citation and access purposes.

Individual scientists can easily share ideas, comments, suggestions and communicate with other members of the community. Trainers can easily share data, algorithms, and technologies and communicate with the classroom. Trainees can easily access to data and technologies in a controlled environment where experiments and tests can be performed, traced and documented.

By exploiting the Social Service any user shall be deemed to accept these conditions of use:

- The sharing process can be done by selecting the privacy level: all the users registered to the Gateway or only the members of a single Virtual Research Environment;
- Ideas the user post and information the user share may be seen and used by other users and D4Science.org cannot guarantee that other Users will not use the ideas and information that any user may share. Therefore, if the user has an idea or information that she/he would like to keep confidential and/or don't want others to use, or that is subject to third party rights that may be infringed by sharing it, that user shall not post it to the Gateway;
- D4Science.org Infrastructure is not responsible for a user's misuse or misappropriation of any content of information any user post on the Gateway Social Service;
- D4Science.org Infrastructure is not responsible on the updates/comments that Users can post on the Social Service;
- In the case the user remove any updates/comments, the D4Science.org Infrastructure cannot guarantee that other Users have copy of the removed content.

4.3.1 Messaging Service

The Social Networking Collaborative Platform implements an Email Service, in order to enable messages exchange only between D4Science.org registered users.

This facility relies on the Social Networking service and on Apache Cassandra for storing, indexing and managing messages.

Any user of the D4Science.org Email Service is subject to the policies reported below. In addition to any other violations described in the D4Science.org terms of use, the user may not:

- Generate or facilitate unsolicited commercial emails;
- Sending unsolicited emails to significant numbers of VRE users belonging to individuals and/or entities with whom the user have no pre-existing relationship;

- Send, upload, distribute or disseminate or offer to do the same with respect to any unlawful, defamatory, harassing, abusive, fraudulent, infringing, obscene, or otherwise objectionable content;
- Transmit content that may be harmful to minors;
- Illegally transmit another's intellectual property or other proprietary information without such owner's or licensor's permission;
- Promote or encourage illegal activity;
- Use the Email Service in connection with illegal peer-to-peer file sharing.

4.3.2 Notification Service

For the purpose of facilitating the user interaction, the Social Networking Collaborative Platform includes a Notification Service. The Notification Service can be configured from the Gateway to an email address associated with the user account.

By exploiting the Notification Service the user shall be deemed to accept these conditions of use:

- The user agrees that the Gateway may receive communication by e-mail;
- The user can review the Notifications settings to enable/disable the email messages notifications;
- The user acknowledges and agrees that D4Science.org shall have no liability associated with or arising from failure to do so, including, but not limited to, user failures to receive important notifications.

4.4 Catalogue Service

The Publication Platform is a comprehensive environment to support data harmonization and publication. It resembles a catalogue of artefacts with search and browse, yet the openness with respect to the typologies of products published, the metadata to document them as well as the integration with the rest make it a flexible environment. Every published item in the catalogue is characterised by (i) a type, which highlights its features and allows an easier search, (ii) an open ended set of metadata which carefully describe the item, and (iii) optional resource(s) representing the actual payload of the item.

This platform primarily relies on CKAN technology, i.e. an open source software enabling users to build and operate open data catalogues. This core technology has been wrapped and extended by means of the Catalogue Service, a component realizing the business logic of the publishing platform. The Catalogue Service enacts the management of Catalogue Items. A Catalogue Portlet, accessible in each VRE, allows navigating the catalogue content as well as publishing content by exploiting the Publishing Widget. This widget is also embedded into the Workspace portlet, so users can publish folders and/or files directly from there. External services can access the catalogue content and publish new items via the gCube Catalogue RESTful APIs. The Catalogue Service relies on the Workspace and Storage Hub for storing the payload of the published items.

Each catalogue item type carefully defines the metadata elements characterising the item typology by specifying the names of the attributes, the possible values, whether an attribute is a single instance or a repeatable one. In addition to that, each item type contains directives on how to exploit attributes for items organisational purposes, e.g. automatically transform values in tags or exploit the values for creating collections or groups of items. Moreover, the following properties apply:

- every catalogue item is equipped with an actionable, persistent, unique identifier (namely a PURL) that can be used for citation and access purposes;
- whenever a catalogue item is published, the associated payload(s) is stored in a persistent storage area to guarantee its long-term availability;

- every catalogue item is equipped with a license carefully characterising the possible (re-)uses;
- every publication of an item leads to the automatic production of a post in the social networking area of the VRE to inform its members;
- every catalogue item is equipped with rich and open metadata, i.e. it is possible to carefully customise the typologies of products and the accompanying metadata to the community needs;
- catalogue contents (item's metadata and resources) are made available for consumption by clients by the RESTful API as well as by other standard APIs, e.g. DCAT and OAI-PMH.

Data Managers extend current practices with controlled and standard data publication with the aim to enlarge the access to data products while maintained full ownership and control over the sharing of results. Individual scientists collaborate via standard formats and protocols with selected colleagues since the conception, definition, validation, and sharing of a result. Working groups normalize the exchange of data on activities, enlarge the audience of products, user and data management for the working group.

D4Science.org offers data sharing at different levels: within a Catalogue accessible only in the VRE, accessible in different VREs and publicly accessible via the Gateway service.

By exploiting the Catalogue Service the user shall be deemed to accept these conditions of use:

- Any user can post content, data and/or publication, to D4Science.org only if (a) the user created and own the rights to the content or the user has the owner's express permission to post the content; and (b) the content does not infringe any other person's or entity's rights (including the copyrights, trademarks, or privacy rights) or violate any applicable laws, the Terms of Use, and the Privacy Policy. The D4Science.org Infrastructure can remove content for any infringing reason;
- Any user is responsible for any content the user shares and the consequences of sharing or publishing such content with others or the general public. This includes, for example, any personal information, any confidential data or any unofficial data. D4Science.org Infrastructure is not responsible for the consequences of sharing or posting any personal or other information on its services;
- When any user uses a service to load content into a VRE, the user retains the irrevocable, exclusive, worldwide right and license to use, reproduce, modify, display, remix, redistribute, create derivative works, and syndicate your content in any medium and through any service. Only where the user has specifically shared, or used a specialized service for sharing data, the data become part of a derivative work to be licensed under a schema selected by the owner;
- Each shared set of data is associated to a copyright license. A CC BY-SA license if the dataset is obtained with a mash up process from different sources (a true derivative product), any other license (not technically supported by D4Science.org) in case of unaltered product shared with other Users/VREs. In both cases the metadata must be filled (if not automatically compiled) and a preferred citation must be indicated;
- When any user uses a Service that provides the capabilities to share, transform, readapt, modify, or combine user content with other content, the user grants D4Science.org an irrevocable, non-exclusive, royalty free, perpetual, worldwide right and license to [use, reproduce, modify, display, remix, perform, distribute, redistribute, adapt, promote, create derivative works, and syndicate] the content [in any medium and through any form of technology or distribution] and to permit any derivative works to be licensed under these same license terms;
- Any user accessing shared data must behave according to the copyright license of the dataset. However, the D4Science.org cannot control or audit access to and / or audit use of shared data. It can therefore not impose a shared data policy, but can encourage sharing according to best practices;

- Exposing data within the D4Science.org (i.e. sharing through one or more VREs) will make them accessible only to authorized users, and cannot be accessed from outside.
- The VRE authorized Users can publish data and make it available to the public exposing to general view with no restrictions except for a registration procedure. Each published set of data is associated to a copyright license;
- Published data can be accessed by the general public through the Gateway and through dedicated web applications (e.g. web services);
- Data accessed through the D4Science.org Gateway do not, and should not, include controls over its end use. However, as the data owner or authoritative source for the data, the submitting Organizations or Institutions must retain version control of datasets accessed;
- Once the data have been downloaded from the Gateway, the D4Science.org Infrastructure cannot vouch for their quality and timeliness;
- The D4Science.org Infrastructure cannot vouch for any analyses conducted with data retrieved from the Gateway;
- Derivative work should only be produced by complying with the terms and conditions established in the license of the used dataset(s);
- The derivative work is owned by the legitimate copyright holder identified among the contributors co-owning the new product and is the contact point for any use, re-use/re-distribute request;
- Derivative work may contain data retrieved from the D4Science.org Infrastructure;
- Every dataset should include in its metadata at least a preferred citation, a reference to the infrastructure source dataset and its generation date, and the date that data were accessed or retrieved from the Gateway or any of its services. If the dataset was a composition of multiple datasets with different citations, the owner of the dataset should have added that to the relevant metadata;
- If such derivative work is produced through VREs, D4Science.org generates a default citation, which matches references to sources used for the derivative work and their respective citations;
- Any user must clearly state “D4Science.org cannot vouch for the data or analyses derived from these data after the data have been retrieved from the D4Science.org Infrastructure”.

4.5 Data Analytics Platform

The Data Analytics Platform is a comprehensive and collaborative environment to perform statistical data analysis. It resembles a stand-alone analytics platform, e.g. Weka, with a collection of ready-to-use algorithms and methods, yet it relies on a distributed and heterogeneous computing infrastructure enacting the execution of complex tasks.

The Data Analytics Platform is composed by four main software components: the DataMiner Master, the DataMiner Worker, the Algorithm Importer, and the Algorithm Publisher:

- The DataMiner Master is a web service that orchestrates (distributes, monitors, collects) the processes to the DataMiner Worker(s). Occasionally, the DataMiner Master can execute non-partitionable algorithms locally. The Master is conceived to work in a cluster of replica services operating behind a proxy acting as load balancer. It is offered by a standard web-based protocol, i.e. OGC WPS.
- The DataMiner Worker is a web service in charge of executing the processes it is assigned to by a Master. The service is conceived to work in a cluster of replica services and is offered by a standard web-based protocol, i.e. OGC WPS.
- The Algorithm Importer is a Web Application enabling users to inject new algorithms into the platform by using various programming languages as R, Java, Python, PHP, and more. Injected new algorithms are private, meaning that only the user can access and execute them. That user can easily either share the code and the algorithm with her/his colleagues or make the new algorithm executable by all users of the VRE.

- The Algorithm Publisher is a web service in charge of deploying new algorithms on the cluster of DataMiner Workers.

Non-partitionable algorithms are directly executed on a DataMiner Master instance and possibly use parallel processing on several cores and a large amount of memory. Distributed algorithms use distributed computing with a Map-Reduce approach.

Both Master and Worker services are conceived to be deployed and operated by relying on various providers, e.g. Master and Worker instances can be deployed on private or public cloud providers

The Data Analytics Platform is not a typical Cloud computational engine since it adds to the capabilities to optimally distribute processes to cluster of workers the following features that effectively enables open science:

- Every process hosted by the platform is equipped with an actionable unique identifier that can be used for citation and access purposes;
- the offering and publication of user provided processes (e.g. scripts, compiled programs) by an as-a-Service standard-based approach (processes are described and exposed by the OGC WPS);
- the ability to manage and support processes produced by using several programming languages (e.g. R, Java, Fortran, Python);
- the automatic production of a detailed provenance record for every analytics task executed by the platform, i.e. the overall input/output data, parameters, and metadata that would allow to reproduce and repeat the task are stored into the workspace and documented by a PROV-O-based accompanying record;
- integration with the shared workspace to implement collaborative experimental spaces, e.g. users can easily share datasets, methods, code;
- extensibility of the platform to quasi-transparently rely on and adapt to a distributed, heterogeneous and elastically provided array of workers to execute the processing tasks.

Software Developers can easily test, validate, and run algorithms developed in a large spectrum of software languages on a distributed infrastructure that appears to them as a single coherent system effectively hiding the different technologies and security frameworks. IT Managers can offer a scalable (on demand), controlled (by maximum quota of exploitable resources per user/per day), and secure environment where executing reproducible data analytical algorithms. Data Managers select the appropriate data analytical models among the wide spectrum of models offered with the aim of offering an easy-to-use data dashboard. Individual scientists have access to a wide spectrum of available data analytical models and can easily test, validate, and run algorithms developed in a large spectrum of software languages. Working groups can configure and enrich the data dashboard with the aim of supporting the activities of a working group while hiding the complexity of a powerful, controlled, and secure e-infrastructure that can scale according to the evolving needs of the group.

By exploiting the Data Analytics Platform the user shall be deemed to accept these conditions of use:

- The user may use the Data Analytics Platform to execute processes that are available in a VRE;
- The user can register a new method to the Data Analytics Platform only if (a) the user created and own the rights to the software method or the user has the owner's express permission to post the software method; and (b) the software method does not infringe any other person's or entity's rights (including the copyrights, trademarks, or privacy rights) or violate any applicable laws, the D4Science Terms of Use, Privacy Policy, or any other posted policies;
- The ownership of any intellectual property rights is not in any way transferred to the D4Science.org Infrastructure;

- In the case the user decides to enable other Users to execute it by making her/his software method executable in the VRE, the user remains responsible for any misuse of the software method by other Users;
- At any time, the user can unshare the access to her software method code. All requests are immediately accepted and operated by the D4Science.org Infrastructure;
- At any time, the user can issue a D4Science Software Maintenance Ticket to unregister and remove access to her executable code. All requests are accepted and operated by D4Science.org within 2 working days;
- The D4Science.org Infrastructure can reproduce, modify, and generate derivative works (such as those resulting from translations, adaptations or other changes required by the management capabilities of the D4Science.org Infrastructure);
- The D4Science.org Infrastructure will not otherwise use, reuse software methods for any purpose, except when required to do so by law;
- The D4Science.org Infrastructure is not responsible of the software methods uploaded and hosted by the infrastructure;
- The D4Science.org Infrastructure will not be responsible of any issue regulating intellectual property rights infringement or illegal use of user's software methods;
- The D4Science.org Infrastructure will make reasonable efforts to insure that software methods are persisted. In the rare event of hardware or software Failures caused by failures to a hard drive or power supply, the D4Science.org Infrastructure will make reasonable attempts to restore any registered software method. No guarantee whatsoever is provided on the success of any user's software method recovery;
- If the D4Science.org Infrastructure owner reasonably believes any of the software method violates the law, infringes or misappropriates the rights of any third party, the user will be notified that such software methods be removed from the Workspace or access to it be disabled. If the user does not remove or disable access to the Prohibited software method within ONE business day of notice, the D4Science.org infrastructure manager may remove or disable access to the software method or block access to it;
- The user can use the software method other users share with her/him in any circumstance and for all usage, exploit it in workflows, communicate to the public, including the right to reproduce or display the software method or copies thereof to the public and perform publicly, as the case may be, the software method;
- The user cannot in any circumstance and for any usage redistribute the software method other users share with her/him or copies thereof, lend and rent those software method or copies thereof, sub-license rights in the software method or copies thereof.

4.6 Tabular Data Manager (TDM) Platform

The Tabular Data Manager Platform provides users with a collaborative toolset to collate, harmonize and manage tabular data (e.g. fishery time series, code lists); to produce datasets relevant to the ecosystem approach to fisheries and the conservation of marine living resources. The harmonization includes conversion from private formats to international formats such as SDMX and access to environment where it is possible to use R-programming.

The Tabular Data Manager Platform is composed by three main software components:

- **Tabular Data Portlet:** human centric web application. Allows management of the Tabular Data Service system on a per-user basis, allowing invocation of tabular data service methods and additional functionality with external apps.
- **Tabular Data Service:** main component of the tabular data architecture. It exposes several remote interfaces covering different areas of functionality.
 - **Operation orchestrator:** The operation orchestrator is a component that receives incoming call requests from the service interface and unwinds them into a sequence of operation call. The orchestrator may enforce policies and command automatic operations/validations according to its configuration.

- **Operation modules:** Operation modules are classes that brings functionality to the tabular data service. This functionality can be reached directly with invocations on the Operation Interface. Operation modules can work directly with data on the Data back-end or leverage the cube manager in order to create/clone tables or modify table metadata.
- **Cube Manager:** The cube manager is the lowest level component of the service stack. Its main responsibilities are managing the creation/modification of tables (and their metadata) and acting as a registry for all the created tables, allowing retrieval of tables metadata.
- **Data/Metadata back-end:** This is where raw tables data and metadata are stored and where the service keeps its management data. It relies on PostgreSQL database, opportunity federated in cluster to ensure scalability and availability.

IT Managers are provided with a scalable and controlled environment for teams working on data quality control, data harmonization, and data analysis and mining. Data Managers can exploit a scalable and controlled environment where data provenance and traceability are integrated features of the data dashboard. Individual scientists can import, validate, and harmonize data in order to make them suitable for analytical methods accessible through D4Science or locally. Working groups can support large and distributed teams requiring a common trusted environment ensuring traceability of the operations, roll-back, and supervision before publication.

By exploiting the Tabular Data Manager Service, the user shall be deemed to accept these conditions of use:

The user may use the TDM to store, retrieve, query, serve, and execute content that is owned, licensed or lawfully obtained by you;

The user can post content to the TDM only if (a) the user created and own the rights to the content or the user has the owner's express permission to post the content; and (b) the content does not infringe any other person's or entity's rights (including the copyrights, trademarks, or privacy rights) or violate any applicable laws, the D4Science Terms of Use, Privacy Policy, or any other posted policies;

The ownership of any intellectual property rights is not in any way transferred to the D4Science.org Infrastructure;

In the case the user decides to share it with other Users, the user remains responsible for any misuse of the data by other Users;

In the case the user decides to publish the content into the SDMX Registry provided by D4Science.org, the user remains responsible for any misuse of the data by other Users;

At any time the user can decide to unregister her data. All requests are immediately accepted and operated by the D4Science.org Infrastructure;

The D4Science.org Infrastructure can reproduce, modify, and generate derivative works (such as those resulting from translations, adaptations or other changes required by the management capabilities of the D4Science.org Infrastructure);

The D4Science.org Infrastructure will not otherwise move or distribute your data for any purpose, except when required to do so by law;

The D4Science.org Infrastructure is not responsible of the data uploaded and hosted by the infrastructure;

The D4Science.org Infrastructure will not be responsible of any issue regulating intellectual property rights infringement or illegal use of user's data;

The D4Science.org Infrastructure will make reasonable efforts to insure that data are persisted. In the event of hardware or software Failures caused by failures to a hard drive or

power supply, the D4Science.org Infrastructure will make reasonable attempts to restore your data. No guarantee whatsoever is provided on the success of any user's data recovery;

If the D4Science.org Infrastructure reasonably believes any of the Content violates the law, infringes or misappropriates the rights of any third party, you will be notified and may request that such content be removed from the Workspace or access to it be disabled. If you do not remove or disable access to the Prohibited Content within ONE business day of our notice, the D4Science.org infrastructure manager may remove or disable access to the Prohibited Content or suspend its access;

The user can use the data other users share with her/him in any circumstance and for all usage, reproduce the data, modify the data, and make derivative data based upon the original data, communicate to the public, including the right to reproduce or display the data or copies thereof to the public and perform publicly, as the case may be, the data;

The user cannot in any circumstance and for any usage redistribute the data other users share with her/him or copies thereof, lend and rent those data or copies thereof, sub-license rights in the data or copies thereof.

4.7 Spatial Data Infrastructure (SDI)

The Spatial Data Infrastructure provides users with the capability to store, discover, access, and manage vectoral and raster georeferenced datasets.

The SDI exploits the following technologies: GeoServer equipped with PostgreSQL and PostGIS, GeoNetwork, Thredds. All the exploited technologies are deployed to ensure fault-tolerance and load-balancing. Moreover, they are designed to scale to the limit established and reported in this SLA. All datasets can be described using the standard ISO 19115, that defines what information should exist in the metadata document, and represented in ISO 19139 that defines how metadata conforming to ISO 19115 have be stored in XML format.

D4Science.org ensures that the SDI services are integrated with the other services. In particular, it is worth to report the following:

- It is possible to publish content in Thredds by simply uploading that content in special folders provided by the Workspace service.
- It is possible to generate and/or simply validate the ISO 19139 metadata before publishing them into GeoNetwork. It is possible also to validate the compliancy with the INSPIRE directive if requested.
- It is possible to use any viewer compliant with the OGC standards, as the MapViewer.
- It is possible to use PgAdmin as tool to connect to the Postgresql database exploited by the GeoServer instances by issuing an explicitly personal request that has to be approved by the VRE Manager and enforced by D4Science.org.

Software Developers can exploit powerful OGC-based spatial management services. IT Managers can find a coherent OGC complaint SDI for data storage, access and processing. Data Managers can exploit a comprehensive overview of spatial info and processes. Individual scientists have a spatial data on demand and processing capabilities at their fingerprints. Working groups can develop a comprehensive understanding of the socio-economic state of and pressure of a specific spatial area.

By exploiting the SDI, the user shall be deemed to accept these conditions of use:

Any user can post metadata and/or data, hereafter generically identified as content, to the SDI only if (a) the user created and own the rights to the content or the user has the owner's express permission to post the content; and (b) the content does not infringe any other person's or entity's rights (including the copyrights, trademarks, or privacy rights) or violate any applicable laws, the Terms of Use, and the Privacy Policy. The D4Science.org Infrastructure can remove content for any infringing reason;

Any user is responsible for any content the user shares and the consequences of sharing or publishing such content with others or the general public. This includes, for example, any personal information, any confidential data or any unofficial data. D4Science.org Infrastructure is not responsible for the consequences of sharing or posting any personal or other information on its services;

When any user uses a service to load content into a VRE, the user retains the irrevocable, exclusive, worldwide right and license to use, reproduce, modify, display, remix, re-distribute, create derivative works, and syndicate your content in any medium and through any service. Only where the user has specifically shared, or used a specialized service for sharing data, the data become part of a derivative work to be licensed under a schema selected by the owner;

Each shared set of data is associated to a copyright license. A CC BY-SA license if the dataset is obtained with a mash up process from different sources (a true derivative product), any other license (not technically supported by D4Science.org) in case of unaltered product shared with other Users/VREs. In both cases the metadata must be filled (if not automatically compiled) and a preferred citation must be indicated;

When any user uses a Service that provides the capabilities to share, transform, readapt, modify, or combine user content with other content, the user grants D4Science.org an irrevocable, non-exclusive, royalty free, perpetual, worldwide right and license to [use, reproduce, modify, display, remix, perform, distribute, redistribute, adapt, promote, create derivative works, and syndicate] the content [in any medium and through any form of technology or distribution] and to permit any derivative works to be licensed under these same license terms;

Any user accessing shared data must behave according to the copyright license of the dataset. However, the SDI cannot control or audit access to and / or audit use of shared data. It can therefore not impose a shared data policy, but can encourage sharing according to best practices;

Exposing data within the SDI will make them accessible only to authorized users and cannot be accessed from outside by default. An explicit action has to be performed to make the data accessible from anonymous users.

The VRE authorized Users can publish data and make it available to the public exposing to general view with no restrictions except for a registration procedure. Each published set of data is associated to a copyright license;

Published data can be accessed by the general public through the Gateway and through dedicated web applications (e.g. web services);

Data accessed through the D4Science.org Gateway do not, and should not, include controls over its end use. However, as the data owner or authoritative source for the data, the submitting Organizations or Institutions must retain version control of datasets accessed;

Once the data have been downloaded from the Gateway, the D4Science.org Infrastructure cannot vouch for their quality and timeliness;

The D4Science.org Infrastructure cannot vouch for any analyses conducted with data retrieved from the Gateway;

Derivative work should only be produced by complying with the terms and conditions established in the license of the used dataset(s);

The derivative work is owned by the legitimate copyright holder identified among the contributors co-owning the new product and is the contact point for any use, re-use/re-distribute request;

Derivative work may contain data retrieved from the D4Science.org Infrastructure;

Every metadata of a dataset should include at least a preferred citation, a reference to the infrastructure source dataset and its generation date, and the date that data were accessed or retrieved from the Gateway or any of its services. If the dataset was a composition of multiple datasets with different citations, the owner of the dataset should have added that to the relevant metadata;

Any user must clearly state "D4Science.org cannot vouch for the data or analyses derived from these data after the data have been retrieved from the D4Science.org Infrastructure".

4.8 Gateway Service

The Gateway Engine service provides users upon registration with the capability to access one or more VREs.

The Gateway service employs the Liferay portal as the portlet-hosting platform. Liferay Portal offers a complete platform for building web apps, mobile apps, and web services quickly, using features and frameworks designed for rapid development, good performance, and ease of use. It runs on all major application servers and servlets containers and it is JSR 168 and JSR 286 compliant. Portlets are using JSR 286 and several other technologies such as Java Server Pages for dynamically generation of HTML/XML documents in response to a client's request, and the most popular Front-End frameworks such as Angular, React or Vue.

With the registration procedure the user accepts to deposit her/his Personal Profile Data and the user shall be deemed to accept these conditions of use:

Among the types of Personal Profile Data that the Gateway collects, by itself or through third parties, there are: Cookies, Usage data, first name, last name and email address.

The Personal Profile Data may be freely provided by the User, or, in case of Usage Data, collected automatically when using an Application;

All Data requested by the Gateway is mandatory and failure to provide this Data may make it impossible for this Gateway to provide its services. In cases where this Gateway specifically states that some Data is not mandatory, Users are free not to communicate this Data without any consequences on the availability or the functioning of the service;

Users who are uncertain about which Personal Data is mandatory are welcome to contact the Provider;

Any use of Cookies – or of other tracking tools – by this Gateway or by the owners of third-party services used by this Gateway serves the purpose of providing the service required by the User, in addition to any other purposes described in the present document and in the Cookie Policy, if available;

Users are responsible for any third-party Personal Data obtained, published or shared through this Gateway and confirm that they have the third party's consent to provide the Data to the Owner;

D4Science.org, i.e. the Data Controller, processes the Data of Users in a proper manner and shall take appropriate security measures to prevent unauthorized access, disclosure, modification, or unauthorized destruction of the Data;

The Data processing is carried out using computers and/or IT enabled tools, following organizational procedures and modes strictly related to the purposes indicated. In addition to D4Science.org, in some cases, the Data may be accessible to certain types of persons in charge, involved with the operation of the site (administration, legal, system administration) or external parties (such as third-party technical service providers, email carriers, hosting providers, communications agencies) appointed, if necessary, as Data Processors by D4Science.org. The updated list of these parties may be requested from D4Science.org at any time;

The Data is processed at the D4Science.org operating offices and in any other places where the parties involved with the processing are located. For further information, please contact D4Science.org;

The Data is kept for the time necessary to provide the service requested by the User, or stated by the purposes outlined in this document, and the User can always request that D4Science.org suspend or remove the data;

The Data concerning the User is collected to allow D4Science.org to provide its services, as well as for the following purposes: Interaction with external social networks and platforms, Contacting the User and Analytics.

DRAFT

5 VRE SPECIFIC SERVICES TERMS OF USE

5.1 Share Latex (fully integrated)

For the purpose of facilitating the operation of the OPENMOD VRE, the **ShareLaTeX** technology will be fully integrated. ShareLaTeX is an online LaTeX editor that allows real-time collaboration and online compiling of projects to PDF format.

In comparison to other LaTeX editors, ShareLaTeX is a server-based application, which is accessed through a web browser. On July 20, 2017, ShareLaTeX was acquired by Overleaf. Overleaf plans to continue ShareLaTeX under the brand Overleaf v2 which was in beta testing up until the 4th of September 2018.

The provision of the ShareLaTeX service is therefore limited to the availability of the open source license and it will be guaranteed by D4Science.org until Overleaf will not change it.

By exploiting the ShareLaTeX, the customer shall be deemed to accept these conditions of use:

The customer agrees that the infrastructure may log, audit, and reboot the service;

The customer acknowledges and agrees that D4Science.org shall have no liability associated with or arising from failure to provide the service, including, but not limited to, failures to emit availability notifications.

By providing the ShareLaTeX, D4Science.org shall be deemed to accept the following conditions of use:

The ownership of any intellectual property rights generated by exploiting this service is not in any way transferred to the D4Science.org Infrastructure;

At any time, the customer can decide to unpublish content made available through ShareLaTeX. All requests are immediately accepted and operated by the D4Science.org Infrastructure;

If the D4Science.org Infrastructure reasonably believes any of the content made available through ShareLaTeX violates the law, infringes or misappropriates the rights of any third party, the customer will be notified and may request that such content be removed from the infrastructure. If the customer does not remove or disable access to the prohibited content within ONE business day of the notice, the D4Science.org infrastructure manager may remove or disable access to that content or suspend its access.

6 SERVICE CAPACITIES

The following service capacities are integral part of this Service Level Agreement.

VRE Management

- 1 VRE with capacity to grow up to 20 VREs upon fifteen (15) working days prior written notice to the Provider;
- Each VRE can be managed with tailored policies to access dedicated storage space and computing capacities; All VREs contribute to the capacities reported for each single service.

Workspace Service

- 0,5 TB with capacity to grow to 2 TB upon fifteen (15) working days prior written notice to the Provider;
- The storage capacity is counted by summing all the capacity used by the users and it includes all the data either directly uploaded by the user or generated through the exploitation of any of the service reported in this Service Level Agreement.
- Ability to deliver 'safe spaces' for non-public data;
- Ability to organize the data storage in several layers (directories/files);
- Ability to manage access rights per directory;
- Ability to assign privileges to users to enable uploading and downloading of files and directories;
- Ability to link distributed legacy data from third party servers accessible through standard protocols.

Messaging Service

- Unlimited usage (within terms and conditions).

Social Networking Service

- Unlimited usage (within terms and conditions).

Notification Service

- Unlimited usage.

Catalogue Service

- Unlimited usage within terms and conditions;
- The data and documents published through the Catalogue service contribute to the workspace capacity as reported in this Service Level Agreement.

Tabular Data Manager

- Unlimited usage within terms and conditions;
- The data and documents published through the Tabular Data Manager service contribute to the workspace capacity as reported in this Service Level Agreement.

Data Analytics Platform (Optional)

- Ability to process at least 15 executing jobs at a time with capacity to grow to 50 executing jobs at a time upon five (5) working days prior written notice to the Provider;
- GUI-based single run and batch processing capabilities (job creation, parameter specification and data selection);
- Ability to port existing software (Java, Python, Octave);
- Ability to manage job submission, monitoring, re-execution;
- Ability to visualize the result of a job;
- Ability to download the result of a job;
- Ability to execute a single set of code in 16 parallel CPU threads for each of the 15 executing jobs at a time.

Gateway Service

- Ability to register 300 users with capacity to grow to 1000 users upon five (5) working days prior written notice to the Provider.

Spatial Data Infrastructure

- Ability to manage vectoral and raster georeferenced datasets through the following set of services: 1 GeoServer (version 2.10.5) with capacity to grow to 3 GeoServer for enabling low latency and load distribution, Varnish server for each GeoServer for enabling caching proxy, 1 GeoServer Load Balancer, 1 Postgresql databases equipped with PostGIS with capacity to grow to 2 Postgresql for enabling replication, 1 Thredds server (version 4.6.9) with capacity to grow to 2 Thredds servers for enabling low latency and load distribution, 1 GeoNetwork (version 3.2.1).
- The data and metadata published through the SDI contribute to the workspace capacity as reported in this Service Level Agreement

ShareLaTex

- Ability to manage 20 articles at a time with capacity to grow to 50 articles upon five (5) working days prior written notice to the Provider.

7 SERVICES HOURS AND EXCEPTIONS

The Services operate during the following hours: twenty-four (24) hours a day, seven (7) days a week, three hundred sixty-five (365) days a year.

The following exceptions apply:

- Planned maintenance windows or service interruptions (“scheduled downtimes”) will be notified via e-mail at least 72 hours before the start of the outage;
- Three scheduled downtimes per year are envisaged to ensure security updates;
- Each scheduled downtime cannot exceed 4 working hours;
- Downtime periods exceeding 6 hours need justification;
- Emergency downtime to answer to critical security alerts.

DRAFT

8 SERVICE LEVEL TARGETS

The table below summarizes agreed service provisioning aspects regarding availability, reliability, and quality of support.

Table 1. Service Level Targets

Service provisioning feature	Target	Notes
Availability	95%	This is defined as the ability of a service or service component to fulfil its intended function at a specific time or over a calendar month. The target value represents the minimum percentage per month.
Reliability	98%	This is defined as the ability of a service or service component to fulfil its intended function at a specific time or over a calendar month, excluding scheduled maintenance periods. The target value represents the minimum percentage per month.
Quality of Support	Medium	See Section 10.

9 SERVICE LEVEL COMPENSATION

This section reports the cost compensation

- for the provision of the Services listed in Sec. 3 (The Services);
- under the condition and policies reported in Sec. 4 (The D4science services terms of use);
- and the customer designed conditions and policies reported in Sec. 5 (VRE specific services Terms of Use);
- with the capacities reported in Sec. 6 (Service Capacities);
- and operated within the limit reported in Sec. 7 (Services hours and exceptions);
- and with the quality of service reported in Sec. 8 (Service level targets);
- and with the Medium support level as reported in Sec. 10 (Support).

The Total Service Cost includes the costs for providing support during the validation phase, the deployment and configuration costs for the setup of the environment and of the services, the training modules. A cost per month with a minimal commitment of one (1) Year is also indicated.

Table 2 presents the supported combinations of services (bundle). The bundles range from the minimal set of services released with minimum capacities (Silver) to a fully-fledged configuration including all the services equipped with the maximum capacities (Platinum). A tailored configuration of the services including key features for a production quality environment is also reported and recommended (Gold).

The Deployment and Configuration Cost is to deliver the OPENMOD VRE and all its services as reported in this SLA.

The first version will be evaluated and the feedback will be used to deliver a second version. The evaluation period depends on the number of services to assess as follow.

- 20 working days for the Silver bundle
- 25 working days for the Gold bundle
- 30 working days for the Platinum bundle

The validation phase will be assisted by a senior engineer that will illustrate and support the Costumer in performing the validation.

The second version will be evaluated by the Costumer in 5 working days independently by the chosen bundle.

Bundle		Silver	Gold	Platinum			
Deployment and Configuration	VRE Management, Messaging, Notification, Social Networking	€5.000	10 VRE Management, Messaging, Notification, Social Networking	€6.000	50 VRE Management, Messaging, Notification, Social Networking	€10.000	
	Workspace (0.5 TB)	€1.000	Workspace (1 TB)	€2.000	Workspace (2 TB)	€3.000	
	Gateway (up to 300 users)	€2.000	Logo, Template, Tailored domain (up to 500 users)	€5.000	Fault-tolerant, low latency, load balancing (up to 1000 users)	€6.000	
	Catalogue	€1.750	Tailored metadata types, Tailored grouping	€2.750	Tailored vocabularies	€4.250	
	Tabular Data Manager				Tailored code-lists	€4.250	
	Method Engine		15 executing jobs at a time	€3.500	50 executing jobs at a time	€5.500	
	Spatial Data Infrastructure	GeoServer				Fault-tolerant, low latency, load balancing	€5.750
		GeoNetwork					€1.000
		Thredds				Fault-tolerant, low latency, load balancing	€5.750
	Validation Cost	€2.000		€4.000		€5.000	
Total Cost		€11.750		€23.250		€50.500	
Training	VRE Management and Operation Training Module	€750		€750		€750	
	Method Engine			€1.500	Method Engine and SDI Training Module	€2.500	
	Workspace and Catalogue Service Training Module	€750		€750		€750	
	Total Cost		€1.500		€3.000		€4.000
Total Service Cost		€13.250		€26.250		€54.500	
Cost per month		€920		€1.636		€2.455	
Cost for the first year		€24.295		€45.886		€83.955	
Cost for the second (up to the fifth) year		€11.045		€19.636		€29.455	
Total cost for 5 years operation		€68.477		€124.432		€201.773	

Table 2. Compensation Costs per Bundle

The Cost per Month is to operate and maintain the VRE and to provide support upon request. It has been calculated as a reimbursement of the personnel cost needed to maintain the service operational, estimated in 5 working days per month (Silver configuration) of a skilled senior engineer costing 4,500 Euro per month. The operation and maintenance working days were calculated according to the following considerations.

Maintenance and operation costs are typically classified into routine and periodic costs. Routine maintenance costs involve small tasks that are undertaken frequently. These tasks are typically cyclical preventive maintenance and inspections tasks that exploit the Accounting

Dashboard to analyse the load and the exploitation of the resources. In the occurrence of peaks of the load, additional computational resources are automatically added but the infrastructure personnel is triggered and they are called to inspect that their functional behavior is correct. Periodic maintenance costs relate to works undertaken at intervals of either several weeks or months. These periodic tasks are usually related to security fixes and hardware upgrade.

Type of Costs (in working days)	Year Estimation			Monthly average		
	Silver	Gold	Platinum	Silver	Gold	Platinum
Time-related costs						
Routine Maintenance	12	24	42	1	2	3,5
Periodic Maintenance	18	36	48	1,5	3	4
Usage-related costs						
Routine Maintenance	24	36	54	2	3	4,5
Total Costs	54	96	144	4,5	8	12

Table 3. Maintenance and Operation Effort in Working Days per Bundle

This obligation is independent of and is not contingent upon the Customer execution of a purchase order. Payment shall be paid, in advance, to the Provider in instalments as set forth in the following schedule of amounts and due dates:

- Upon 15 days from the Service Level Agreement signature: 30% of the Total Service Cost;
- Upon 15 days from the execution of the Service Level Agreement with the release the environment and of all services: 70% of the Total Service Cost;
- The cost per month will be accumulated for a period from the 1st day of the first calendar month after the provision of the environment to the last day of the following third calendar month ("Accumulation Period"). The Customer will initiate payment for invoices collected during the Accumulation Period on the first business day of the month nearest to fifteen (15) days following the end of the Accumulation Period.

Payment will be in Euro currency unless otherwise stated. Payment will not constitute acceptance of deliverables.

9.1 Training Module Cost

The following training modules are envisaged.

Module	Location	Material	Trainee	Trainers	Cost
VRE Management and Operation Half day training module covering the following services: VRE Management Service, Messaging Service, Social Networking Service, Notification Service, Gateway Service	<i>Provider premises</i>	<i>Slides</i>	20	2	750 Eur.
Workspace Service Half day training module covering the Workspace Service	<i>Provider premises</i>	<i>Slides</i>	20	1	375 Eur.

Catalogue Service Half day training module covering the Catalogue Service	<i>Provider premises</i>	<i>Slides</i>	20	1	375 <i>Eur.</i>
Data Analytics Platform and SDI One full day training module covering the Data Analytics Platform	<i>Provider premises</i>	<i>Slides</i>	20	2	1500 <i>Eur.</i>
Data Analytics Platform and SDI Two full day training module covering both the Data Analytics Platform and the Spatial Data Infrastructure	<i>Provider premises</i>	<i>Slides</i>	20	2	2500 <i>Eur.</i>

The Costumer can host training modules upon compensation of the travel costs as following:

A daily subsistence allowance of 80 € shall be granted where the distance between the Provider premises and the place specified in the Costumer's letter of invitation is more than 50 kilometres. It shall only be paid in respect of the day(s) on which the training module is executed plus two additional days (corresponding to the day before the training event and the day after the training event) where the distance between the Provider premises and the place specified in the Costumer's letter of invitation is more than 200 kilometres;

An accommodation allowances up to EUR 150 per night shall be granted where the distance between the place of residence and the place specified in the candidate's letter of invitation is more than 200 km. This allowance shall be granted in cases where the candidate has to stay one or two more nights because of incompatibility between the times specified in the letter of invitation and transport timetables. It shall be paid on presentation of an accommodation invoice confirming the timeframe of the stay and payment;

A travel allowance corresponding to either the train or air fare (on presentation of the either train or flight ticket invoice).

10 SUPPORT

Support is provided via D4Science.org Service Desk accessible via the Gateway service. Access requires a valid D4Science token or the login via the Gateway. Support is available between:

- Monday to Friday;
- From 9:00 to 18:00 in the time zone of the relevant Resource Centre.

Service times always apply with the exception of public holidays in the country of the supporting Resource Centre.

10.1 Incident handling

Incidents will be handled according to the Quality of Support level that is estimated according to the impact of the outage or service quality degradation.

The Quality of Support in this Agreement has level: Medium, so the incidents, based on their priority will be responded to with the following response times:

Table 4. Response times to incidents according to the incident priority of "Medium" services table

Incident priority	Response time	Comment
Less urgent	10 working days	wishes and enhancements that are "nice to have"
Urgent	5 working days	service available but performance degraded; no work-around available
Very Urgent	2 working days	service available but performance degraded; work-around available
Top Priority	1 working day	service interrupted;

10.2 Service requests

In addition to resolving incidents, standard service requests (e.g. change requests, information requests, documentation) will be fulfilled through the defined support channels in the same way as incidents. Service requests are classified as "Less urgent".

10.3 Service Operation

All D4Science services are provisioned in accordance with the D4Science Terms of Use that defines the rules by which any user must agree to abide in order to use the service. Moreover, the ways D4Science gathers, uses, discloses, and manages client's personal information are reported in the Privacy Policy. Finally, all the services are operated according to well-defined service operations procedures as reported in Table 5.

Table 5. Service Operation Procedures and Rules

Rules	Specification
Terms of Use	https://services.d4science.org/terms-of-use
Privacy Policy	https://www.iubenda.com/privacy-policy/441050
Decommissioning Procedure	https://wiki.d4science.org/index.php?title=Procedure_Decommissioning_Procedure
Disaster Recovery Plan	https://wiki.d4science.org/index.php?title=Procedure_Disaster_Recovery_Plan

Metrics And Statistics	https://wiki.d4science.org/index.php?title=Procedure_Infrastructure_Accounting_Dashboard
Operations Documentation	https://wiki.d4science.org/index.php?title=D4Science_Deployment_and_Operation
Service Accounting	https://wiki.d4science.org/index.php?title=Procedure_Infrastructure_Accounting
Service Availability Hours	https://wiki.d4science.org/index.php?title=Procedure_Service_Availability_Hours
Service Level 1 Support	https://wiki.d4science.org/index.php?title=Procedure_Infrastructure_Incident_Management#1st_Line_Support
Service Level 2 Support	https://wiki.d4science.org/index.php?title=Procedure_Infrastructure_Incident_Management#2nd_Line_Support
Service Level 3 Support	https://wiki.d4science.org/index.php?title=Procedure_Infrastructure_Incident_Management#Management_Decision
Service Level Support Hours	https://wiki.d4science.org/index.php?title=Procedure_Infrastructure_Incident_Management#Prioritization_and_Assignment
Service Maintenance Window	https://wiki.d4science.org/index.php?title=Procedure_Downtime_Declaration
Service Monitoring	https://wiki.d4science.org/index.php?title=Procedure_Infrastructure_Monitoring

11 LIMITATIONS AND CONSTRAINTS

The Services provisioning is subject to the following limitations and constraints.

- Support is provided in English;
- Availability and Reliability calculations are based on the D4Science Service Monitoring operational results;
- Failures in the Service Monitoring are not considered as SLA violations;
- Downtimes needed to ensure the security of the Service issues are not considered Agreement violations;
- Failures of resource provider not being part of D4Science production infrastructure are not considered as Agreement violations;
- Force Majeure. A party shall not be liable for any failure of or delay in the performance of this Agreement for the period that such failure or delay is due to causes beyond its reasonable control. Means any
 - fire, flood, earthquake or natural phenomena;
 - war, embargo, riot, civil disorder, rebellion, revolution.

12 COMMUNICATION, REPORTING AND ESCALATION

12.1 General communication

The following contacts will be generally used for communications related to the service in the scope of this Agreement.

Table 6. Contacts

Customer contact	TBD
Provider contact	Pasquale Pagano pasquale.pagano@isti.cnr.it D4Science Technical Director
Service Support contact	See Section 8

12.2 Regular reporting

As part of the fulfilment of this Agreement and provisioning of the service, the reports in Table 7 will be provided.

Table 7. Reporting

Report title	Contents	Frequency	Delivery
Services Performance Report	The document provides the overall assessment of service performance (per month and per VRE) and SLA target performance achieved during last 4 months	Every four months	Email to the Customer
Platform Service Report	The document provides usage records (resources used, users) of the Customer service during last 4 months	Every four months	Email to the Customer

All reports shall follow predefined templates.

The Provider commits to inform the Customer, if this Agreement is violated or violation is anticipated. The following rules are agreed for communication in the event of violation:

- In case of violations of the Services targets for two consecutive months, the Provider will provide justifications and a plan for Services enhancement to the Customer. The Provider will produce a status report and a Service enhancement plan for the improvement of the Services within one month from the date of the first notification;
- The Customer will notify the supporting Resource Centres in case of suspected violation via the D4Science Service Desk (Section 8). The case will be analysed to identify the cause and verify the violation.

12.3 Escalation and complaints

For escalation and complaints, the Provider contact point shall be used, and the following rules apply:

- In case of repeated violation of the Services targets for two consecutive months, a review of the Agreement and of the Services enhancement plan will take place involving the parties of the Agreement;
- Complaints or concerns about the Services provided should be directed to the Provider contact who will promptly address these concerns. Should the Customer still feel dissatisfied, about either the result of the response or the behaviour of the supporting Resource Centre, the D4Science Director should be informed.

13 INFORMATION SECURITY AND DATA PROTECTION

The following rules for information security and data protection related to the Service apply:

- Assertion of absolute security in IT systems is impossible. The Provider is making every effort to maximize security level of users' data and minimize possible harm in the event of an incident;
- The Provider will define and abide by an information security and data protection policy related to the service being provided;
- The parties of the Agreement will meet all requirements of any relevant policy or procedure of the Provider and will comply with the applicable national legislations;
- The detailed information security and data protection are reported in the VRE Terms of Use document. It is integral part of this SLA and is tailored for the VRE specific needs. However, they cannot in any case and for any reason violate the D4Science Terms of Use and the D4Science Privacy Policy.

14 RESPONSIBILITIES

Neither party will use the name of the other party, or any name of the other party's employee(s), in any publicity, advertising, or news release without the prior written approval of an authorized representative of that party. The Customer shall, however, acknowledge the Provider's support for the investigations being pursued under this Service Level Agreement. In any such statements, the relationship of the parties shall be accurately and appropriately described.

Neither party is authorized or empowered to act as agent for the other for any purpose and shall not on behalf of the other enter into any contract, warranty, or representation as to any matter. Neither shall be bound by the acts or conduct of the other.

During the Term of the Service Level Agreement, the parties may be disclosing Proprietary Information to each other. Prior to disclosure by the Customer, the Customer shall inform the Provider of its intent to disclose Proprietary Information; the Customer and the Provider shall have the right to decline receipt of said information. Each party agrees to treat Proprietary Information received from the other with the same degree of care with which it treats its own Proprietary Information, but not less than reasonable care, and further agrees not to disclose such Proprietary Information to a third party without prior written consent from the disclosing party.

The foregoing obligations of non-disclosure do not apply to Proprietary Information which:

- Was known to the recipient prior to the disclosure hereunder;
- Was received by the recipient without restriction from a third party who was not under an obligation of confidence to the Provider, the Customer or a third party regarding such information;
- Is in the public domain at the time of disclosure hereunder or subsequently enters the public domain without the fault of the recipient;
- Has been independently developed by personnel of recipient without access to such Proprietary Information, and recipient can substantiate any claim of independent development by competent proof; or
- Is required to be disclosed pursuant to a valid subpoena, court order or other requirement of applicable law. Prior to making any such disclosure the disclosing party shall make all reasonable efforts to advise the other party of the required disclosure so that the other party may seek legal protection for such information.

Unless otherwise agreed to in writing, the obligations under this section shall terminate one (1) year after the date of receipt of Proprietary Information.

14.1 Of the Provider

Additional responsibilities of the Provider are as follows:

- The Provider adheres to all applicable operational and security policies and procedures and to other policy documents referenced therein;
- The Provider allows monitoring of the Service in order to measure the fulfilment of the agreed service level targets;
- The Provider retains the right to introduce changes in how the Service is provided, in which case the Provider will promptly inform the Customer and update the Agreement accordingly.

Limitation of responsibilities of the Provider:

- The Provider shall not be responsible or liable with respect to any subject matter of this Service Level Agreement under any contract, negligence, strict liability or other theory for any indirect, incidental, special or consequential damages including, but not limited to, loss of revenues and loss of profits;

- The Provider's liability to the Costumer for any claim related to or arising from this Service Level Agreement or the relationship created by this Service Level Agreement shall be limited to the amount paid by the Costumer to the Provider pursuant to Section 9 Service Level Compensation;
- The Provider will operate the services in accordance with generally accepted professional standards. The Provider makes no representations or warranties, express or implied, and disclaims any implied warranties, including, without limitation, the warranties of merchantability, fitness for a particular purpose or infringement of intellectual property.

14.2 Of the Customer

Additional responsibilities of the Customer are as follows:

- The Customer commits to acknowledge D4Science in the scientific publications benefiting from the Service: "This work used the D4Science Infrastructure with the support of CNR-ISTI";
- The Customer will provide during Agreement review (yearly) list of scientific publications benefiting from the Service;
- The Customer must not share access credentials with anyone else;
- The data stored in the system by the Customer must not cause any legal violation due to the content type (such as copyright infringement, dual use, illegal material);
- The use must be consistent with the Terms of Use of the Service;
- The Customer will notify the Provider in case the actual amount of the Service used results in being under- or over-estimated. The Customer will request an update of the Agreement to ensure optimal usage of the Service;
- The Customer will create one or more Groups to define the group of users entitled to access the Service;
- The Customer must request D4Science Service Desk support to enable assigning tickets with appropriate Group name;
- The Costumer shall defend, fully indemnify and hold harmless the Provider, its trustees, officers, employees, students, agents, successors, heirs and assigns against all claims arising out of the Costumer's use, commercialization, or distribution of information, materials, products or reports which result in whole or in part from the activities performed pursuant to this Service Level Agreement, including without limitation claims for personal injury, property damage, and infringement of intellectual property.

Limitation of responsibilities of the Costumer:

- The Costumer will in no event be liable for any direct or indirect, material or moral, damages of any kind, arising out of the Service Level Agreement or of the use of the Services, including without limitation, damages for loss of goodwill, work stoppage, computer failure or malfunction, loss of data or any commercial damage, even if the Costumer has been advised of the possibility of such damage.

15 REVIEW, EXTENSIONS AND TERMINATION

The Services performance will be reviewed against the defined Service level targets according to Section 6. The Agreement will be annually reviewed until expiration.

If the Customer wishes to extend the duration after the Agreement termination date, an extension will be negotiated with the Provider.

The Provider retains the right to introduce changes in the Service, in which case the Customer retains the right of terminating the Agreement.

The Agreement can be terminated at any time upon agreement of the parties.

The foregoing obligations of termination are extended under the following conditions:

- Either party may terminate this Agreement upon ninety (90) days prior written notice to the other after the first year of operation counted from the first day of this Service Level Agreement;
- In the event that either party hereto commits any material breach of any terms or conditions of this Service Level Agreement, and also fails to reasonably remedy such breach within sixty (60) days after receipt of written notice thereof, the non-breaching party may, at its option and in addition to any other remedies that it may have at law or in equity, terminate this Service Level Agreement by sending notice of termination in writing to the other party to such effect. Termination shall be effective as of the day of the receipt of such notice;
- Termination of this Service Level Agreement by either party for any reason shall not affect the rights and obligations of the parties accrued prior to the effective date of termination of this Service Level Agreement.

Amendments, comments and suggestions must be addressed to the Provider and the Customer according to Section 12.

16 GENERAL PROVISIONS

16.1 Non-assignability

The rights and obligations of the parties under this Service Level Agreement shall not be assigned without the prior written permission of the other party.

16.2 Entire Agreement

This Service Level Agreement contains the entire and only agreement between the parties respecting the subject matter hereof and supersedes or cancels all previous and contemporaneous negotiations, agreements, commitments and writings between the parties on the subject of this Service Level Agreement, including, but not limited to, non-disclosure agreements between the Provider and the Customer. Should processing of this Service Level Agreement require issuance of a purchase order or other contractual document, all terms and conditions of said document are hereby deleted in their entirety. This Service Level Agreement may not be amended in any manner except by an instrument in writing signed by the duly authorized representatives of each of the parties hereto.

16.3 Governing Law

Any litigation arising between Parties and resulting from the interpretation of this Service Level Agreement will be subject to the exclusive jurisdiction of the competent court where the Provider resides or conducts its primary business.

16.4 Force Majeure

Neither party shall be responsible to the other for failure to perform any of the obligations imposed by this Service Level Agreement, provided such failure is occasioned by fire, flood, explosion, lightning, windstorm, earthquake, subsidence of soil, failure or destruction, in whole or in part, of machinery or equipment or failure of supply of materials, discontinuity in the supply of power, governmental interference, civil commotion, riot, war, strikes, labor disturbance, transportation difficulties, labor shortage or any other cause beyond its reasonable control.