Tool Augmentation by user enhancements and Orchestration

Service Level Agreement for TAO Services

Title	Service Level Agreement for TAO Services
Reference	CSRO/S2AS/CIU/CV/26/0052
Issue	1.0
Date	2026-02-06
Status	Final

This page has been intentionally left blank.

Change record

Date	Issue	Section	Page	Comment
2025-08-11	1.0	All	All	Initial issue
2026-04-27	2.0	5	19	Added disclaimer for data sources.

Control document

Process	Name	Date
Written by:	Ramona Popescu	2025-08-11
Checked by:	Cosmin Udroiu	2026-04-27
Approved by:	Cosmin Udroiu	2026-04-27

Table of Contents

1. Introduction

1.1. Purpose and Scope

This Service Level Agreement ("Agreement") is entered into between CS GROUP – ROMANIA S.A. ("Operator") and the User, to govern the delivery of the TAO services ("Service"). The relevant contacts and representatives are identified in section 5.

The purpose of this Agreement is to establish clear expectations and commitments between the Operator and the User for the reliable and effective performance of the Service. This Agreement shall be valid for two years from the date of acceptance and may be extended beyond its expiration.

1.2. Document Structure

This section describes the structure of the document.

- Section 1 is this introduction
- Section 2 defines the parties, role and responsibilities related to the services
- Section 3 describes the service components
- Section 4 defines the Service objectives and measurements including the service level targets and reporting means
- Section 5 identifies the communication channels

1.3. Acronyms and Abbreviations

Term	Definition
Agreement	A document that describes a formal understanding between two or more parties.
Availability	Ability of an IT service or other configuration item to perform its agreed function when required.
Escalation	An activity that obtains additional resources when these are needed to meet service level targets or customer expectations.
Event Management	The process responsible for managing events throughout their lifecycle.
Failure	Loss of ability to operate to specification, or to deliver the required output.
First-Line Support	The first level in a hierarchy of support groups involved in the resolution of incidents.
Impact	A measure of the effect of an incident, problem or change on business processes.
Incident	An unplanned interruption to an IT service or reduction in the quality of an IT service.
Incident Management	The process responsible for managing the lifecycle of all incidents.
Platform Tier Service Provider	The individual economic operator, or prime contractor with sub-contractors, or Consortium lead (or lead of comparable legal structures) owning the Platform Tier Service and acting as the NoR tender submitter
Priority	A category used to identify the relative importance of an incident, problem or change.
Process	A structured set of activities designed to accomplish a specific objective.
Request Fulfilment	The process responsible for managing the lifecycle of all service requests.
Resolution	Action taken to repair the root cause of an incident or problem, or to implement a workaround.
Role	A set of responsibilities, activities and authorities assigned to a person or team.
Service Hours	An agreed time period when a particular service should be available.

Service Level	Measured and reported achievement against one or more service level targets.
Service Level Agreement (SLA)	An agreement between a service provider and a customer.
Service Level Target	A commitment that is documented in a service level agreement.
Service Request	A formal request from a user for something to be provided
Supplier	An economic operator entered into legal commitment with the prime operator of the Platform (the Platform Tier Service Provider) in order to deliver resources (e.g. data, processing services) via the Platform
Urgency	A measure of how long it will be until an incident, problem or change has a significant impact on the business.

2. Parties, Roles and Responsibilities

2.1. Parties

This SLA is between the user and the TAO Operator:

The Service Provider (TAO Operator)
CS GROUP – ROMANIA S.R.L. Str. Pacii nr. 29 Craiova, Dolj, ROMANIA +40 (0)251 41 28 50 https://www.c-s.ro/

Table 1: TAO Operator Contacts

2.2. Contacts

The following persons have negotiated this document and agree it will be used as the formal Service Level Agreement (SLA) for the provision of the Service.

Name	Position
For "The Service Provider"	
Sorin SCORTAN	Key Accounts Manager, CS GROUP – ROMANIA S.R.L.
Cosmin Udroiu	Platform Operations Manager, CS GROUP – ROMANIA S.R.L.

Table 2: Contacts

2.3. The Service Provider responsibilities

The Service Provider will provide and maintain all the service(s) indicated in the section Service Description.

Additionally, the Service Provider will:

- Ensure relevant services are available to the User in line with the uptime levels listed in the next sections.
- Respond to relevant user support requests within the timescales listed below.
- Take steps to investigate and resolve issues within the timescales listed below.
- Maintain good communication with the users at all times.
- Maintain good communication with the Customer at all times.

2.4. The Customer responsibilities

The Customer will:

- Notify the Service Provider of issues or problems in a timely manner.
- Provide the Service Provider with all information available about reported issues or problems to facilitate investigation and resolution
- Provide prompt feedback on the delivered service access grants to allow coordination of Platform user on-boarding and measurement of Platform user adoption rates
- Maintain good communication with the Service Provider at all times.

3. Service Description

This SLA applies to the following services (that will be collectively indicated as "The Service" in the following) provided by the TAO Platform:

General TAO Introduction

The TAO platform, developed in the frame of an ESA project, is a lightweight, consistent, and easy to use orchestration framework, easily scalable to hundreds or thousands processing nodes. It allows the reuse of commonly used processing toolboxes such as (but not limited to) Orfeo Toolbox and SNAP/Sentinel Toolboxes and their enhancement with capabilities for processing composition and distribution in such a way that end-users can define by themselves workflows, without in-depth IT or programming knowledge, and easily integrating custom processing modules (either third party processing modules or written by users using popular scripting languages).

TAO framework allows for processing composition and distribution in such a way that end users could define by themselves processing workflows and easily integrate additional processing modules (by processing module it is understood either a standalone executable or a script).

In terms of use, the TAO platform provides a mean for orchestration of heterogeneous processing components and libraries in order to process scientific data. This is achieved in following steps:

- Preparation of resources (including processing components) and data input
- Definition of a workflow as a processing chain
- Execution of workflows
- Retrieval / visualization of the results

A processing component represents a standalone application defined by the following parameters:

- Input description: type of data that the component accepts as input source (e.g. image, raster maps, vector maps, sensors, etc.)
- Processing operation with execution parameters: the operation that the component executes with the list of accepted parameters.
- Output description: type of data provided as processing operation result. Can consist in one or more files.

A component is viewed as a generic execution resource and should permit, through its parameters list, the possibility to define *input location* (the place where it expects to retrieve input data for consuming) and *output location* (the place where the processing results are persisted).

The processing modules are heterogeneous regarding their implementation language or operating system. TAO aims to be as OS-independent as possible and to allow the integration of modules written in different programming languages (such as C/C++, Java or Python).

TAO allows for workflows creation by combining individual components, named processing components. They are heterogeneous regarding their implementation language or operating

system and can be written in different programming languages (such as C/C++, Java or Python). Along with out-of-the-box components from well-known EO toolboxes (OTB, SNAP, GDAL), custom code components can be added by users. These components can be edited in a Jupyter Notebook, hence several programming languages are supported, given the variety of language kernels for Jupyter. Currently, the platform supports writing components in Python, R, Ruby and Javascript.

By workflow it is understood a directed acyclic graph consisting of processing operations performed on a given input (EO data and/or ancillary data), having at least one output. A workflow is described by a formal flow diagramming technique, with directed flows between components.

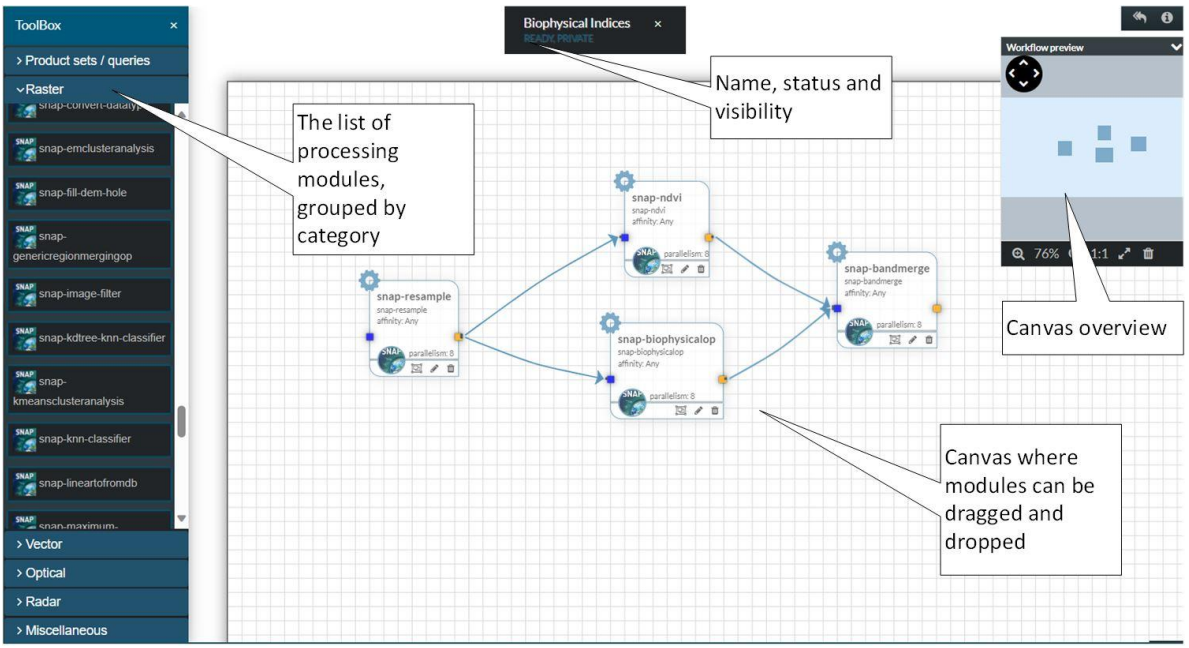

The modules are grouped into categories that represent packaged applications.

A workflow can be created from scratch, can be edited, and deleted. It can also be created (or cloned) from an existing workflow. Users can find workflows based on some tags or filter existing ones.

As it is OGC Application Packaging compatible, TAO also provide the possibility to the user to import workflows and components from existing CWL workflows. The system will automatically create all the necessary Docker images descriptors, the Processing Components descriptors and the internal TAO workflow corresponding to the imported CWL workflow. Vice-versa, once the user has built its algorithms as TAO workflows, he can export them into CWL, Argo, Bash and Javascript workflows that can be executed or imported on other platforms.

In terms of functionality, it is possible to query and select the products on which the algorithm is applied as well as to specify some of the algorithm parameters. To support this kind of services TAO offers a set of data sources that can be used for querying and fetching EO products (download or directly use the products when available locally, ex. in /eodata). Users

can create sites defined by their AOI and temporal extent that can be used for simpler later queries to the different data sources available. From the already downloaded EO products user can create, by having the possibility of filtering the products, the so called product datasets that can be used as inputs to different processing services available in TAO. Alternatively, the data sources can be directly connected to the inputs of the different processing services in which case, the products are downloaded automatically according to the querying parameters, and the services workflows are fed with all the products returned by the query. The user has also the possibility to upload additional user data like LPIS or similar structured data or other files that can be used as inputs in the algorithms. Results are stored online in a user private space and can be downloaded on the user's machine. It is possible to share the processing results with other people on the platform or publicly on the internet. In his/her account, the user can also mount different external repositories and in this moment several types are supported:

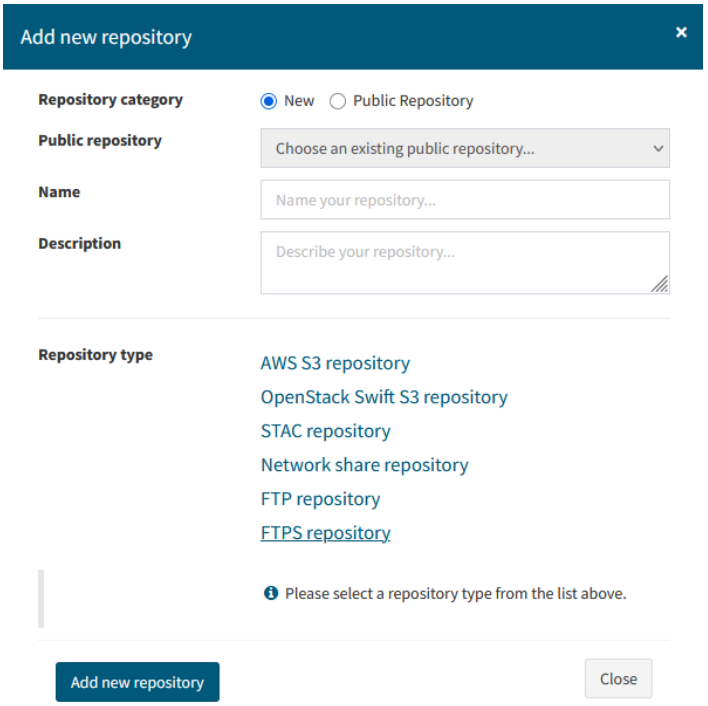

It is worth noting that TAO already includes several public repositories.

Into his repositories administration page, the user can copy easily data from the remote repositories to the local repository or vice versa (if allowed).

- STAC repositories
- S3 repositories
- FTP and network repositories.

In terms of execution of the defined user workflows, TAO implements several Resource Managers that implement DRMAA interface (Distributed Resource Management Application API) allowing local execution, SSH executions, Kubernetes execution and dynamic allocation and deallocation of virtual machines in cloud environments where OpenStack is available. Using dynamic allocation of virtual machines and deallocation them when no longer used help

in optimizing the resources cloud costs while also ensuring a vertical and horizontal scalability of the executions by both parallelizing the executions but also choosing the right flavour for the tools executing.

TAO implements a dual-mode subscription mechanism to manage user access to computational resources, built on top of OpenStack, under two main models: fixed resources and pay-per-use.

- Pay-Per-Use: users are charged based on actual resource consumption, particularly the number and type of virtual machines (VMs) instantiated. Resource usage is tracked via a monitoring system, ensuring billing reflects real-time usage, although the precision of this tracking may vary.
- Fixed Resources: users operate under a predefined subscription flavor, which specifies the number and type of VMs allocated. When a user attempts to create or terminate a VM, the system checks this operation against their subscription flavor. The request is only authorized if sufficient resources are available under the user's plan, ensuring compliance with the allocated quota.

Data Processing Services

This service offers access to a set of algorithms (i.e. processing services) which can be applied to one or more EO products to obtain a specific information.

The Platform is equipped with a large set of EO processing services including:

- Sentinel-1 pre-processing into Amplitude and Coherence products starting from the Sentinel-1 SLC products
- Biophysical and Spectral indicators calculation (LAI, FAPAR, FCover, NDVI, NDWI, Brightness etc.) starting from the Sentinel-2 and Landsat-8 L2A products
- Grassland mowing detection containing information about grassland mowing events over the parcels declared by the farmers as permanent grassland in a given region or country. It is based on the processing of S1 and S2 derived products, which are respectively coherences, calibrated amplitude backscatter and three BIs which are the NDVI, the FAPAR and the LAI.
- L3 Basic Markers aiming to produce products that contain the S1 and S2 signals which are aggregated over each declared parcel (assumed to be homogeneous land management unit), moving from pixel to object analysis. For each parcel, signal statistics and markers are then computed and stored in a collection which is therefore growing over time
- Agricultural Practices monitoring that aims at the identification of the agricultural practice of crop harvesting/clearance and the comparison of the temporal behaviour of vegetation of fallow lands, catch-crops or nitrogen-fixing crops with the pre-defined (country specific) EFA (Ecological Focused Area) practices rules
- Crop Type Mapping detection generates a parcel-level crop type map based on a dense time series of Sentinel-2 and Sentinel-1 acquisitions and on an in-situ dataset, nominally the subsidy applications. It trains a Random Forest model from a subset of

the in-situ dataset, which is then applied on the entire set of parcels. The predicted crop type can then be compared to the expected one, to check the conformity of the declarations.

- Bare soil detection that aims providing products containing bare soil and non-bare soil markers, based on the optical Sentinel-2 (S2) surface reflectance and Synthetic Aperture Radar (SAR) Sentinel-1 (S1) time series at the parcel level.
- Parcel Heterogeneity detection for detection the homogeneity of the parcels in terms of the crop uniformity.
- Change of agricultural category detection that allows detection from one year to another of a change in the agricultural category of a certain parcel
- Tillage monitoring performs the tillage detection for a certain parcel
- Winter Cover Crops that aim to provide products related to the future Winter Cover Crop (WCC) processor related to the WCC detection, WCC season length estimation and WCC aboveground biomass estimation
- Burned area monitoring performs the detection of burned areas.

The Data Processing Services operated on TAO are split in two subcategories according to the complexity of the underlying algorithm as specified in Annex 1 of this EULA document:

- Simple services
- Advanced services

User Algorithm Hosting

This service also based on TAO is meant for developer users to plug onto the Platform already existing scientific applications written in a variety of languages (e.g. Java, C++, Python, R), and then deploy, automate, manage and scale them in a very modular way.

It offers access to a development environment in the Platform Cloud infrastructure and an online resource exposing a REST API that allows the developer user to upload his algorithm (also developed offline) onto the platform and deploy it as a Platform-operated data processing service.

The User Algorithm Hosting Service enables the use of such algorithms with the EO data and user private data and takes care about the access to and use of processing resources. It makes possible the sharing of the algorithm with other users of the platform.

4. Service Objectives and Measurements

4.1. Service Coverage

The Service operates during the following hours:

- The **Data Processing Services**: 24/7, 365 days a year
- The **User Algorithm Hosting:** 24/7, 365 days a year
- The **User Support Services:** during normal working days (NWD) and hours (NWH)

with the following definitions:

Coverage	Definition
NWD	Normal Working Day (Monday-Friday except on Christmas Day, 26th December, New Year's Day, 2nd January, 6th January, 7th January, 24th January, Easter Friday, Easter Monday, 1st May, 1st June, Pentecost, 15th August, 30th November and 1st December)
NWH	Mon-Fri: 9:00-18:00 EET
24/7	24 hours per day, 7 days a week

Table 3: Service coverage definitions

4.2. Service Downtime

Planned downtime for routine maintenance of the Service, upgrades or problem resolution:

- will not be performed during NWH, except for correction of blocking issues
- will be scheduled, as far as possible, with 5 NWDs notice and agreed with the Customer
- users will be informed of any impact on services via email and via a post in the forum with 4 NWDs notice, with information on the functionality involved, the affected users, the nature and the expected duration of service interruption.
- In case of extension of the downtime, the note shall be followed up with an update.

In case of unplanned downtime, the Service Provider will inform the affected users as soon as the problem is detected with the same type of information as for planned downtimes. In addition, the Service Provider will issue a confirmation note of service restoration after every unplanned downtime.

4.3. Service Performance and Reliability

The Service is provided according to the KPIs reported in the following sub-sections.

A **failure** will be considered any Incident with **Criticality = Blocking** where the following definitions apply:

Table 4: Incidents Criticality

Incident Criticality	Definition in scope of Operations
Blocking	A blocking incident in the TAO service occurs when the platform is unable to continue its core orchestration or workflow execution activities, resulting in a complete halt of service for end-users. Such an incident prevents users from preparing resources, defining workflows, executing processing chains, or retrieving results, effectively rendering the platform non-operational. Examples: - **Authentication / Authorization failure** – users cannot log in or their quotas are incorrectly enforced, blocking workflow execution. - **Data access disruption** – local or remote data repositories (e.g., Copernicus SciHub, AWS, USGS) are unreachable or corrupted, making EO product retrieval impossible. - **Cluster execution failure** – the Distributed Resource Manager (e.g., SLURM or Torque via DRMAA) becomes unavailable, preventing workflow jobs from being scheduled or executed. - **Processing container errors** – Docker execution environments fail to initialize or isolate processing modules, leading to critical runtime dependency conflicts. - **Storage outage or quota enforcement malfunction** – local file system or product database becomes unavailable, blocking both input and output data operations.
Critical	A critical incident in the TAO service is a severe disruption that significantly degrades the platform's functionality but does not necessarily block all users or operations. For example: - **Workflow execution degradation** – workflows run with excessive delays due to resource contention, misconfigured quotas, or failing execution nodes in the cluster. - **Partial data access failures** – one or more external repositories (e.g., SciHub, AWS, USGS, DIAS) cannot be queried or return incomplete results, limiting the datasets available for processing. - **Container/runtime instability** – some Docker containers hosting processing modules (e.g., SNAP, OTB, Python, R) fail to initialize, causing only certain workflows to fail while others continue to run. - **Node unavailability** – one or several computing nodes become *offline* in the cluster, reducing overall capacity and causing workflow scheduling delays. - **Monitoring or quota enforcement malfunction** – incorrect reporting of CPU, memory, or storage usage, leading to users being prematurely restricted or consuming excessive resources that affect others.
Routine	All other issues that do not affect the delivery of the service to users

4.3.1. Data Processing Services

KPI-ID	Description
DPS_1	**Node Availability (%)**: > 99 % uptime per month
DPS_2	**Workflow Success Rate (%)**: ≥ 95% (target, excluding user misconfigurations)
DPS_3	**Average Data Retrieval Time:** - Local repository: < 60 seconds. - Remote repositories (Copernicus, AWS, USGS): 10–20 minutes per product
DPS_4	**Average Task Execution Time:** - Simple services: 3–20 minutes. - Advanced: variable, typically several hours, depending on AOI and season interval

Table 5: KPI for Data Processing Services

Simple services

KPI-ID	Description
DPSA-1	Algorithm execution service availability is at least 99% on a monthly basis
DPSA-2	Time between the launch and start of execution of an algorithm is less or equal than 5 sec
DPSA-3	Less than 5% of algorithm executions fails due to platform / algorithm issues

Table 5: KPI for conventional Data Processing Services

Advanced services

KPI-ID	Description
DPSC-1	Algorithm execution service availability is at least 99% on a monthly basis
DPSC-2	Time between the launch and start of execution of an algorithm is less or equal than 5 sec
DPSC-3	Less than 25% of executions fails due to platform / algorithm issues

Table 6: KPI for advanced Data Processing Services

4.3.2. User Algorithm Hosting

KPI-ID	Description
UAH_1	**Algorithm Onboarding Success Rate (%)** ≥ 95% (uploaded algorithms successfully registered in TAO without missing dependencies)
UAH_2	**Execution Success Rate (%):** ≥ 90% for user algorithms
UAH_3	**User Algorithm Error Reports (#):** < 5 per 100 executions

Table 6: KPI for User Algorithm Hosting

KPI-ID	Description
UAH-1	Algorithm deployment service availability is at least 99% on a monthly basis
UAH-2	Time to make a new algorithm available as a data processing service from the availability of the Application Package is less or equal to 5 NWD
UAH-3	Time to share a deployed algorithm with other users is less or equal than 1 NWD

Table 7: KPIs for User Algorithm Hosting

4.4. Measurement Details and Reporting

The monitoring and measurement of the KPIs defined in this agreement will be performed through the TAO platform's built-in monitoring dashboard and supporting logging mechanisms.

- **Frequency**: Measurements will be collected continuously, with reporting consolidated on a monthly basis.
- **Data Sources**: System logs, workflow execution reports, quota usage records, and monitoring database entries.
- **Responsibility**: The Service Provider is responsible for collecting and validating KPI measurements, while the Customer has read-only access to relevant dashboards and reports.
- **Calculation**:
  - *Node Availability (DPS_1)* – ratio of total node uptime to total time in a month.
  - *Workflow Success Rate (DPS_2)* – percentage of workflows completed without error, excluding user misconfigurations.
  - *Average Data Retrieval Time (DPS_3)* – average time measured between request submission and data delivery from local/remote repositories.
  - *Average Task Execution Time (DPS_4)* – mean duration for tasks grouped by processing module type (e.g., SNAP/OTB vs. user scripts).

- *Algorithm Onboarding Success Rate (UAH_1)* – percentage of user algorithms registered without dependency or packaging errors.
- *Execution Success Rate (UAH_2)* – ratio of successful executions to total executions for user algorithms.
- *User Algorithm Error Reports (UAH_3)* – number of incident/error tickets raised per 100 user algorithm executions.

KPI results will be compared against the agreed thresholds to determine compliance with the SLA.

Event Management will provide the environment for all the measurement and reporting required to monitor the components, the process and the Service itself. The primary source for monitoring the targets in this SLA is the information collected into the Service Support Desk management tool (i.e. JIRA Service Desk). In the case of measuring elapsed times (like resolution time which is pivotal for rating the Service), the value is calculated from the time the incident or request is created in the system until the time the incident or request is marked as closed.

When the incident or service request is reported outside the working hours, the record is actually created at the beginning of the next working hour, unless a special procedure is invoked to start resolution immediately.

5. TAO Data Access and External Sources Disclaimer

TAO provides users with integrated access to multiple external data sources, including but not limited to EO-CAT, USGS, Alaska Satellite Facility, EarthData, and Copernicus Data Space. Access to such external data sources may require users to supply their own valid credentials (e.g., username and password), which are managed and controlled solely by the respective data providers.

TAO does not own, host, or distribute data originating from these external sources. Instead, TAO acts as an intermediary interface that facilitates user-initiated access to and retrieval of data from third-party systems. As such, TAO does not guarantee the availability, integrity, performance, or continuity of access to any external data source. Any limitations, restrictions, outages, or changes imposed by the external providers are outside the control and responsibility of TAO.

In cases where data is made available directly by the underlying cloud infrastructure hosting TAO (for example, through provider-managed repositories such as mounted data volumes), such data may be accessed without the need for external credentials. However, the availability and accessibility of such data remain subject to the policies, service levels, and operational conditions of the respective cloud provider.

Users are solely responsible for:

- Obtaining and maintaining valid credentials for external data sources;

- Complying with the terms of use, licensing conditions, and access policies of each external data provider;
- Ensuring that their use of the data adheres to applicable legal and regulatory requirements.

TAO shall not be held liable for:

- Any failure, delay, or interruption in accessing external data sources;
- Any loss, corruption, or inconsistency of data retrieved from third-party providers;
- Any misuse or unauthorized use of user credentials;
- Any changes in access conditions, pricing, or availability imposed by external data providers or cloud infrastructure operators.

By using TAO to access external data sources, users acknowledge and accept that TAO operates solely as a facilitation layer and that all data access is ultimately governed by the respective external providers and/or cloud infrastructure services.

6. Communication

The following contact shall be used for customer inquiries related to the scope of this SLA:

- Service Provider contact for the Customer: cs.ro.office@soprasteria.com

More specifically, all communications between the Service Provider and the Customer regarding both incidents and service requests will take place through email ensuring traceability, timely updates, and proper escalation procedures.

7. Annex 1: TAO Processing Services

6.1. Simple EO data processing services

Maintainer	Service
CS GROUP - ROMANIA	Sen4CAP S1 Pre-processing
CS GROUP - ROMANIA	Sen4CAP Spectral indices
CS GROUP - ROMANIA	Sen4CAP Biophysical indicators
CS GROUP - ROMANIA	L3 Basic Markers

6.2. Advanced EO data processing services

Maintainer	Service
CS GROUP - ROMANIA	Sen4CAP Crop Type Mapping
CS GROUP - ROMANIA	Sen4CAP Agricultural Practices
CS GROUP - ROMANIA	Sen4CAP Grassland Mowing Detection
CS GROUP - ROMANIA	Sen4CAP Bare soil Detection

Maintainer	Service
CS GROUP - ROMANIA	Sen4CAP Parcel Heterogeneity Detection
CS GROUP - ROMANIA	Sen4CAP Tillage Detection
CS GROUP - ROMANIA	Sen4CAP Winter Cover Crops Detection
CS GROUP - ROMANIA	Burned Area Detection