ECMWF Copernicus Procurement

Invitation to Tender

Copernicus Climate Change Service

An LLM-based data search and discovery service for the Climate Data Store

Volume II: Specification of Requirements

ITT Ref: C3S2_523
ISSUED BY: ECMWF Administration Department Procurement Section
Date: 5 August 2025
Version: Final

Funded by the European Union

Implemented by ECMWF

Table of Contents

1 Introduction

The Copernicus Climate Change Service (C3S) implemented by ECMWF on behalf of the European Union develops and delivers authoritative, quality-assured information about the past, current and future states of the climate in Europe and worldwide. It aims to a) inform policy development to protect citizens from climate-related hazards such as high-impact weather events, b) improve the planning of mitigation and adaptation practices for key human and societal activities, and c) promote the development of new applications and services for the benefit of society.

Evaluation and Quality Control (EQC) is a central component of C3S to establish the service as a trusted source of climate information, delivering quality-assured and authoritative service outputs such as datasets and applications that are traceable and reproducible. The EQC function ensures transparency of the service outputs including their quality attributes and builds the basis for a true operationalisation of climate services and the inclusion of climate data into policies and standards. This transparency builds trust in the data and supports responsible, effective use in both operational and policy contexts.

Information about fitness for purpose of the datasets is essential for users helping them determine whether a dataset or application is suitable for their specific needs, decisions, or applications. C3S users come from diverse backgrounds—ranging from scientists and policymakers to industry professionals and the general public—and have varying levels of expertise. Clear and actionable information about data quality, uncertainty, limitations, and appropriate use ensures that users can confidently assess relevance, avoid misinterpretation, and make reasonable and science-based decisions.

Recent advances in AI and machine learning, particularly in natural language processing (NLP) and large language models (LLM), have made it possible to develop tools that can understand and generate human-like text, interpret complex queries, and extract relevant insights from vast datasets. These models can now comprehend domain-specific language, summarise information, and engage in interactive dialogue, making them ideal for navigating and explaining complex climate and environmental data. This progress enables the creation of intuitive, intelligent interfaces that bridge the gap between expert data and user needs, transforming how users discover, understand, and apply scientific information.

ECMWF as the Entrusted Entity for the Copernicus Climate Change Service invites tenders to develop an advanced LLM-based search and discovery service, to enhance user access to climate and environmental data from the C3S Climate Data Store (CDS). The service will enable natural language querying, intelligent data summarisation, and personalised recommendations to support a broader and more diverse user base, including non-experts. It will integrate with existing systems like the CDS Virtual Assistant (“Knowledge Duck”) and explore the use of generative AI for automated quality assessments of datasets.

The successful Tenderer shall be responsible for the provision of an intuitive, intelligent interface that improves data discoverability, trust, and usability for climate services.

2 Background Information

The purpose of this section is to clarify the context of this tender and to describe the vision for EQC as well as relationships and possible synergies of the envisaged tool with ongoing activities. Specific technical requirements for the work to be carried out under this tender are described in Section 3 of this document.

2.1 Climate Data Store (CDS)

The backbone of the C3S is the cloud-based Climate Data Store (CDS) [https://cds.climate.copernicus.eu/] that provides users with a single point of access to quality-assured climate and meteorology data. The datasets may be stored in different data centres worldwide or in remote servers, but this complexity is transparent to CDS users. C3S data is offered with open access and is free to use under the CC-BY data licence

and other licenses. Data are properly documented and enriched by appropriate quality attributes provided by the EQC function. All CDS data and tools are accessible from webforms as well as via open Application Programming Interfaces (APIs).

The Copernicus Climate Data Store (CDS) has expressed a clear aspiration to move toward full compliance with the FAIR data principles — *Findable, Accessible, Interoperable, and Reusable*. The CDS has already made demonstrable progress toward full FAIR compliance, achieving a 70% FAIR score via independent assessment. (see Zanacchi, M., Damasio da Costa, E., Glinton, S., & Cioni, G. (2025). *Towards a FAIRer Data Stores service*. ECMWF Newsletter No. 183, Spring 2025. Retrieved from https://www.ecmwf.int/en/newsletter/183/news/towards-fairer-data-stores-service).

The CDS data catalogue provides access to climate datasets via a searchable catalogue., such as satellite and in situ observations, reanalyses, seasonal forecasts and climate projections. The current search function is a "pure text-based search" without the use of LLMs or other complex AI techniques.

2.2 Vision for Evaluation and Quality Control

C3S has established an EQC framework for all its products and services to ensure that users are served well and that this will continue to be the case as their needs evolve. The main goal of the EQC function is to develop informative and accurate statements about the quality associated to well-identified use cases. Those statements, in combination with other documented information about the datasets, constitute a knowledge base that can help users to assess fitness for purpose, given their needs and requirements. The vision is that a user can employ a range of tools for locating, selecting, and accessing climate data, simply by expressing requirements in their own language. This vision is explained in a paper by Dee et al. (2024): *"Are Our Climate Data Fit for Your Purpose?"*—published in the *Bulletin of the American Meteorological Society*. It emphasises that while the CDS offers a vast array of climate datasets, the utility of these datasets heavily depends on their suitability for specific applications. A structured approach is proposed for assessing the datasets, focusing on aspects such as accuracy, consistency, and relevance to user needs, recognising the diverse requirements of users, from researchers to policymakers. The paper advocates for evaluation methods that consider the specific purposes for which the data will be used.

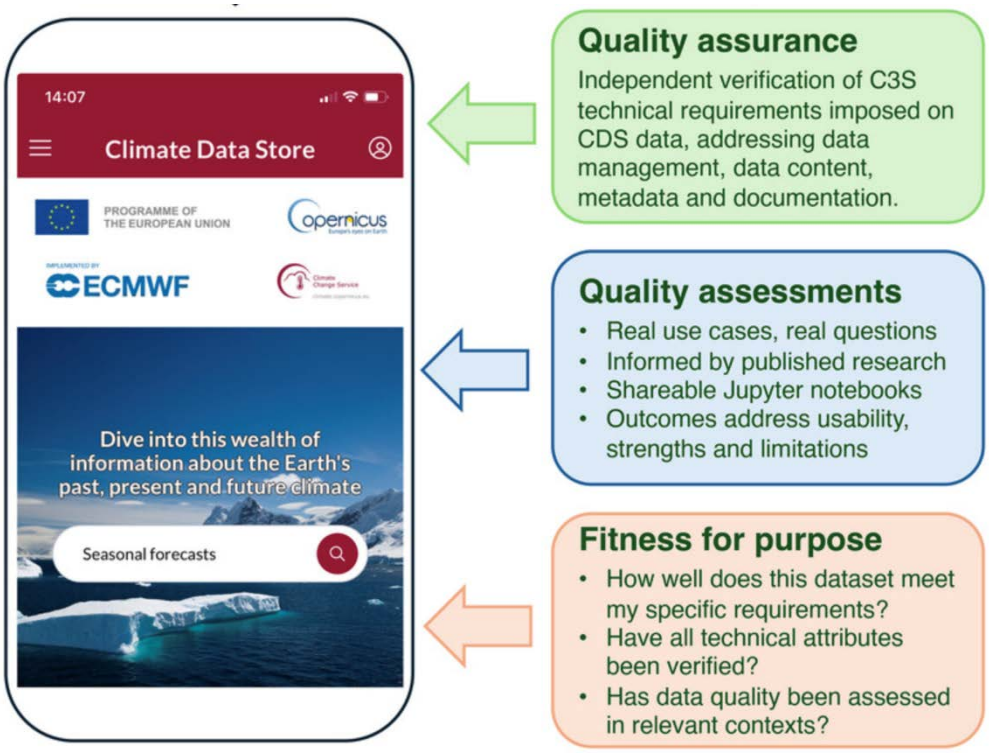

Figure 1: The three pillars of the C3S EQC for Datasets Framework 2.0

The EQC framework makes a distinction between quality assurance and quality assessment. Quality assurance serves to inform users that data, metadata and documentation comply with a well-defined set of verifiable technical requirements. It provides evidence that this compliance has been checked independently from the producers. Quality assurance for each CDS dataset is implemented by verifying a set of well-defined technical requirements associated with the dataset. The purpose of quality assessments, on the other hand, is to provide science-based information about accuracy, uncertainties, sources of uncertainty, temporal consistency, strengths, and weaknesses of a dataset in the context of realistic use cases. In many instances, especially for datasets that are widely used, relevant information is available in the published literature. EQC evaluators are tasked with developing quality assessments that are designed to generate useful statements about fitness for purpose of CDS datasets. The assessments address concrete questions about data quality associated with real use cases. Many of the assessments are implemented in Jupyter notebooks that can be shared, re-used, and modified by users. Assessments can involve multiple CDS datasets and use other sources of reference data. Quality assessments are offered in a Jupyer Book, which can be accessed here: https://ecmwf-projects.github.io/c3s2-eqc-quality-assessment/intro.html

Taken together, the outcomes of these activities provide the key information needed to determine fitness for purpose (Figure 1).

A team of independent experts continually checks, evaluates and documents the quality of all CDS datasets, to ensure that they meet user requirements and are fit for purpose. Currently, the EQC function is implemented by the contracts C3S2_520 (EQC for datasets) and C3S2_521 (EQC for applications, indicators, derived datasets and learning resources).

2.3 Synergies with existing tools

Tenderers are invited to look at several prototypes and projects that may be relevant for this work including:

ChatECMWF

This tool was initiated during a Code4Earth challenge as an experimental AI-driven assistant designed to help users explore ECMWF's open data and climate services. It leverages large language models to answer natural language queries and guide users toward relevant datasets or documentation. The project demonstrated the potential for conversational interfaces to improve accessibility and engagement with Copernicus data but remains a prototype with limited integration into operational systems. (https://codeforearth.ecmwf.int/project/chatecmwf-2023/)

ClimSight

The paper "Local climate services for all, courtesy of large language models" by Koldunov & Jung (2024) presents ClimSight, a proof-of-concept tool that leverages LLMs to deliver localised, user-friendly climate information. By integrating LLMs with geospatial data and climate model outputs, ClimSight can interpret natural language queries and provide relevant climate insights tailored to specific locations. This approach addresses the challenge of making complex climate data accessible and actionable for a broad audience, including non-experts. In the context of the tender, this paper underscores the potential of LLM-based tools to democratise climate services, aligning with the goal of developing an intuitive search and discovery tool that enhances user engagement with climate and environmental data.

CDS Virtual Assistant

The Climate Data Store Virtual Assistant (VA), also known as the 'Knowledge Duck', is a Google Dialogflow powered tool that enhances user support by providing immediate assistance through a pop-up interface on CDS web pages. It draws information from various sources, including the Knowledge Base, User Forum, ECMWF parameter database, and existing CDS documentation, to answer user queries effectively. The

assistant also offers pre-configured buttons for frequently asked questions on topics like user accounts, ERA5, and the CDS API. The system is continuously refined based on daily user feedback, aiming to improve response quality and user satisfaction. Future developments include expanding information sources, enhancing the detection of complex queries, and potentially integrating with other ECMWF services to provide a more agile and comprehensive user support experience.

The most popular questions collected by the current CDS VA include the following:

- C3S data anomalies
- Where can I find my CDS API key?
- How can I get daily data?
- What is the CDS API?
- CDS API troubleshooting
- How many requests can I submit?
- How to install the CDS API?
- Why do my requests take a long time?
- How can I access C3S data?
- Is the data for free?
- How are ERA5 and ERA5-land variables defined?
- How can I see the queue system status?
- How can I create an account?
- Why are my requests slow?
- Where is the wind speed?

Destination Earth Chatbot

Within contract DE_394 ("Development and Implementation of a Climate and Weather Chatbot for Destination Earth"), a chatbot is being designed to be an advanced, AI-powered assistant that allows users to interact with complex weather and climate data through natural language. It will enable intuitive question-and-answer interactions, making it easier for a wide range of users—including policymakers, researchers, and the general public—to access, explore, and understand data from Destination Earth's digital twins and services. The chatbot will use cutting-edge natural language processing (NLP) and machine learning to interpret queries, guide users to relevant datasets and insights, and provide context-aware explanations. It aims to bridge the gap between technical data platforms and user needs by offering personalised, real-time support and fostering broader use of high-value climate and environmental data, see also https://destine.ecmwf.int/news/build-destination-earths-weather-and-climate-chatbot-with-us/

Comparison with existing off-the-shelf solutions

While general-purpose language models such as ChatGPT and Gemini demonstrate impressive capabilities in natural language understanding, they lack access to authoritative, structured climate data and the quality control frameworks essential for scientific credibility. The tool to be developed in this tender goes beyond generic AI by integrating directly with the Copernicus Climate Data Store (CDS), the Evaluation and Quality Control (EQC) framework, and domain-specific metadata. It delivers scientifically grounded dataset recommendations based on fitness-for-purpose criteria, supports multilingual and role-based user interactions, and explains complex metadata in context. This purpose-built system will bridge the gap between users and climate data through a tailored, trustworthy interface that is fully aligned with Copernicus operational workflows — a functionality that cannot be replicated by standalone commercial LLMs.

3 Technical Requirements

The main objective of this contract is the development of an LLM-based service to enable users search and discover suitable dataset in the context of their application. The service includes a tool, a web interface and other activities specified in this section.

The proposed solution is expected to function as a user-facing service, that shall a) operate with full availability and accessibility via the CDS, b) provide a natural language interface, c) come with user support and comprehensive documentation, d) offers regular maintenance mechanisms, and e) includes capabilities for monitoring and logging of user querying for continuous service improvement.

The service performance shall be monitored with a set of KPIs described further below.

The successful Tenderer shall:

- Analyse user needs and define system requirements for the LLM-based tool, ensuring alignment with diverse user profiles and use cases.
- Design and develop the LLM-based search and discovery tool, including natural language understanding, summarisation, and recommendation functionalities.
- Integrate the tool with existing platforms, such as the Climate Data Store (CDS) and the CDS virtual assistant (Knowledge Duck). The tool is expected to gather information from all sources the current virtual assistant uses and provide response to users' queries. This includes structured information contained in a knowledge base, but also unstructured information saved in past Jira tickets and the ECMWF user forum without personal information. The tool should offer the same functionalities which are already successfully delivered by the current VA, improve its weaknesses and expand beyond its current capabilities.
- Implement and test AI-based quality assessment features as part of an experimental work package to enhance data trust and usability.
- Conduct user testing and validation to evaluate the tool's performance, usability, and fitness for purpose across various user groups.
- Deliver comprehensive documentation, training materials, and recommendations for future development, integration, and maintenance.
- Ensure synergies with the chatbot developed within Destination Earth, aligning architecture, user experience, and knowledge sources to support interoperability and avoid duplication of efforts.

The Tenderer shall demonstrate proven expertise in climate and environmental data, including reanalysis, climate projections, seasonal forecasts, and observational datasets. This expertise must be reflected in the composition of the proposed team, including the presence of senior scientific staff capable of verifying and validating outputs from the LLM tool. Close collaboration with respective technical officers at ECMWF is expected.

Below is a list of the core functions of the LLM-based Search and Discovery Tool:

1. Natural Language Querying

- Allow users to ask questions in plain language (e.g., "*Which dataset should I use for precipitation trends in southern Europe?*") in multiple languages.
- Understand both general and domain-specific climate terminology.

2. Dataset Recommendation Engine

- Suggest appropriate datasets based on the user's query, purpose, and context (e.g., timescale, region, variable).
- Offer reasoning for each recommendation (based on evaluation criteria like accuracy, resolution, update frequency) and point towards relevant scientific literature.

3. Metadata Interpretation

- Translate complex metadata into plain language.
- Highlight key dataset attributes: coverage, temporal resolution, limitations, known issues, intended use cases.

4. Fitness-for-Purpose Evaluation Assistant

- Implement the framework discussed in the article: help users assess whether a dataset fits their specific use case.
- Prompt users with relevant questions (e.g., "Do you need near-real-time data?") to refine suggestions

5. User Profiling & Customization

- The tool shall support interaction modes tailored to different user profiles (e.g., policy user, researcher, analyst, general public). These profiles determine the level of detail, complexity of language, and nature of dataset recommendations. In addition to explicit role selection, the system shall include functionality for a posteriori profiling, where users may be offered profile suggestions based on their interaction history. For example: *"Based on your query patterns, you appear to have similar needs to a policy user. Would you like to switch to that profile?"* This adaptive logic improves user experience and supports more accurate response tailoring.

6. Workflow Assistance & Tool Integration

- Guide users through typical workflows (e.g., selecting a dataset → downloading → visualizing).
- Provide users with Python scripts/fragment of code where appropriate, e.g. for advanced users.
- Integrate with existing tools like climatelab / Earthkit and Jupyter Notebooks.

7. Learning & Onboarding

- Offer step-by-step tutorials or quick-start guides.
- Explain climate concepts and how to interpret data (e.g., anomalies, ensemble means, hindcasts).

8. Interactive Dialogue

- Enable follow-up questions or clarifications in a conversational thread.
- Allow users to refine or iterate their queries without starting over.

9. Interface Options

- Available via a web interface (chat UI).

10. User Feedback and Hallucination Logging

- The tool shall include an embedded feedback mechanism (e.g., a "Was this helpful?" or "Report an issue" button) that allows users to flag responses they believe are unclear, misleading, or factually incorrect. These inputs, along with automated logging of potentially hallucinated outputs, shall be systematically reviewed during internal and public testing. Logged issues shall be used to improve model prompts, refine content filtering, and enhance system transparency — supporting continuous improvement and building user trust.

The LLM tool shall cater to a wider range of users, such as decision makers, educators, students, and the general public, without losing its core focus on evaluating data quality and fitness for purpose. This shall be done by enhancing its functionality and adapting its design to suit different levels of expertise and types of user needs.

The following Multi-level User Interfaces shall be included:

- *Beginner Mode:* For the general public and non-experts, the tool will offer simplified, easy-to-understand explanations and visualisations related to datasets and climate concepts, without requiring technical knowledge.
    - Example: A user might ask, *"Which dataset should I use to understand how rainfall has changed in my country over the last 30 years?"* and the tool would respond with a clear explanation, a recommended dataset, and links to relevant visual summaries or guidance material.
- *Intermediate Mode:* For more advanced users the tool can provide more detailed explanations about dataset quality, such as the potential limitations of datasets and their expected accuracy for specific use cases.
    - Example: "This dataset has a spatial resolution of 0.25 degrees, which may be suitable for global trends but not local studies."
- *Advanced Mode:* For expert users, like climate researchers, the tool can provide in-depth technical evaluations and detailed metadata interpretation. This could include guidance on data provenance, biases, statistical analysis, and how those influence the model's predictions.
    - Example: "This reanalysis dataset has a known model bias in the northern hemisphere during winter months."

At the start of the interaction, the tool shall ask a short set of onboarding questions to infer the user's profile (e.g., level of expertise, intended use), and adapt its responses accordingly. This profile may be dynamically reassessed based on the user's query patterns and interaction style, allowing the tool to adjust the level of technical detail or suggest switching to a more suitable interaction mode. This way, the tool can balance complexity with accessibility, ensuring that expert users get precise answers while beginners get clear, jargon-free explanations. The tool shall use adaptive questioning to guide users through the data evaluation process, asking them what their goals are or what they want to achieve. For example: *"Are you looking to make policy decisions based on this data, or do you need it for educational purposes?*" Based on this, the tool could either provide a comprehensive evaluation or a simplified summary. The information on user profile might also be based on the profile stored in the CDS user database.

For non-technical users, the tool shall include visual representations of data quality, such as graphs and charts showing temporal trends, spatial coverage, and resolution. This makes it easier for non-experts to understand the strengths and limitations of different datasets without diving into metadata. Example: A heat map could illustrate areas where a dataset's accuracy might be weaker.

Since many users may not be familiar with scientific jargon, the tool shall offer a built-in glossary or tooltips that explain complex terms in simple language. For instance, it could explain terms like "ensemble mean," "bias correction," or "spatial resolution" in layman's terms when those terms are used in responses. Each response could also come with a summary tailored to the user's expertise. For beginners, the summary could simplify the content, while for experts, it could include technical terms and references to peer-reviewed literature.

The LLM tool shall include built-in tutorials or learning modules that help users at different levels understand how to assess climate data. These resources could guide them through using the tool effectively, whether they're just starting to learn about climate science or are looking for advanced insights into data evaluation.

To make the tool accessible to a wider, international audience, it shall be adapted to support multiple languages, especially for users in regions that are heavily impacted by climate change and where language may be a barrier to understanding.

For users who may have difficulty typing or navigating complex interfaces, the tool shall offer voice interaction and text-to-speech capabilities. This makes it more inclusive for people with disabilities or those in regions with limited access to high-tech devices.

After evaluating datasets for a user's specific needs, the tool shall provide recommendations or reports with actionable insights based on the evaluation criteria (e.g., dataset accuracy, timeliness, resolution). This could be particularly helpful for non-technical users, such as policy advisors who need to quickly understand the most suitable data for a given purpose.

By adapting the LLM tool to cater to a wider range of users, it can become a versatile, multi-functional tool that serves everyone from researchers and policymakers to students and the general public. By simplifying the interface, introducing multi-level modes, adding visualizations, and ensuring inclusivity, the tool can broaden its impact, helping more people make informed decisions about climate data while still retaining its focus on the quality and fitness-for-purpose of that data.

The tool shall not only simplify access to complex climate datasets but provide scientifically grounded, user-specific guidance aligned with Copernicus standards. Unlike generic LLMs and off-the-shelf solutions, its integration with CDS metadata, EQC assessments, and operational climate workflows will provide a level of trust, relevance, and transparency essential to C3S users.

The use of Generative AI to develop quality assessments that feed into the LLM tool shall be explored. This shall be implemented as an experimental work package. By integrating generative AI techniques, the quality assessment process for climate data could be automated and enhanced, providing more sophisticated, scalable, and adaptive assessments that can be dynamically interpreted by the LLM.

The contract shall follow a 30-month implementation schedule, with key delivery milestones including an initial working prototype of the LLM-based tool by Month 10, followed by an internal pre-operational rollout for selected users by Month 15, and a public beta release by Month 21, offering the tool's core functionalities to a broader user group for real-world evaluation and refinement. The remaining period will be used to consolidate user feedback, finalise documentation and training materials, and complete the operational transition and full handover by Month 30.

Milestones and deliverables shall be structured to reflect this timeline, with early user feedback and system refinement occurring in parallel with development.

The proposal shall include indicative cost estimates for the individual content to be produced under each component.

The following subsections list specific technical requirements and tasks of the contractor addressed by this tender.

3.1 WP1: Requirements and System Design

This work package shall define the functional and technical requirements for the LLM-based search tool, including user stories, system architecture, data sources, and integration with existing platforms (e.g., Knowledge Duck, CDS).

Scientific expertise shall be directly involved in the requirements-gathering phase to ensure the tool captures the needs of users requiring high levels of data quality assurance. The climate science lead shall contribute to designing user profiles, defining use cases involving quality evaluation, and ensuring that the architecture accounts for the complexity of metadata and dataset provenance.

The successful Tenderer shall:

- Gather stakeholder input (mainly from ECMWF staff and contractors) and user requirements.
- Define system architecture and integration strategies.

- Ensure the compatibility of the LLM tool with automated data discovery in accordance with the FAIR principles and ensure that the system design supports FAIR compliance within the CDS.
- Detail the LLM tool's core features (e.g., natural language querying, data summarization, personalised recommendations).
- Develop a high-level design for data management, AI model integration, and quality assessment mechanisms.
- Define a roadmap for scalability and performance requirements**.**

Expected deliverables:

- *Stakeholder Requirements Report:* Document detailing the functional and technical requirements based on internal and external input.
- *System Architecture Design:* A comprehensive design document describing the architecture, integration points, content sources and data flows.
- *Hosting Infrastructure:* An estimate of the hosting infrastructure and typology of resources required to deploy and run the system (eg. GPUs).
- *LLM Tool Features Specification:* Detailed specification of LLM tool capabilities, including NLP features, data query handling, and expected performance.
- *Integration Plan:* A clear plan for integrating the LLM tool with Copernicus Data Store (CDS)**,** Knowledge Duck**,** and other relevant systems.
- *Scalability and Performance Requirements:* Document outlining the performance targets, scalability requirements, and resource allocation for the system.
- *Risk Management Plan:* Identification of key risks, mitigation strategies, and assumptions regarding the integration of AI-based quality assessments.

3.2 WP2: LLM Tool Development and Integration

In this work package, the LLM-based search tool shall be developed using an iterative, sprint-based approach, allowing for phased implementation, continuous integration, and regular stakeholder feedback throughout the development cycle.

In this work package, the LLM-based search tool shall be developed, integrating the AI model with existing data platforms and ensuring it provides the required functionality for data search, exploration, and insight generation.

The tool shall be supported by a curated collection of high-quality, domain-relevant sources — including EQC reports, dataset documentation, CDS metadata, and user guidance material — to ensure that its responses are accurate, traceable, and tailored to the Copernicus Climate Change Service (C3S) context. Rather than full-scale model training, these sources shall be indexed or used in retrieval-augmented generation (RAG) workflows, allowing the LLM to ground its answers in trusted Copernicus materials without the need for deep integration or reformatting of internal content.

Sources include the comprehensive quality assessments specifically developed for C3S datasets, available through the C3S Evaluation and Quality Control (EQC) Jupyter Book, which provides detailed insights into the datasets strengths, limitations, and fitness for purpose. In addition, the tool shall incorporate information from selected peer-reviewed literature such as dataset intercomparison studies, scientific publications, and methodological papers that evaluate and benchmark climate data products. These sources will collectively enable the tool to offer informed guidance on dataset selection, interpretation, and appropriate use, supporting a wide range of users in making scientifically sound decisions.

All outputs, including training datasets, quality assessments, and validation procedures, shall be reviewed and verified by domain experts with demonstrated experience in climate data and services to ensure scientific robustness and appropriateness for C3S users.

Since much of this material is dispersed, inconsistently structured, and not yet packaged for use in RAG and hence not yet organised in a format directly usable by LLMs, the successful Tenderer shall identify, extract, and curate a structured collection of authoritative content, including EQC reports, dataset documentation, CDS metadata, user guidance, and learning resources and prepare it for integration into the LLM-based tool. This includes ensuring the information is machine-readable, appropriately formatted for RAG and traceable to its source.

To support early validation and reduce project risk, a prototype version of the LLM-based tool shall be delivered by Month 10. This version will include core NLP functionality, initial CDS metadata integration, and a basic user interface, enabling structured internal feedback prior to the full integration phase.

The successful Tenderer shall:

- Design and implement the LLM model with advanced NLP capabilities for understanding complex, multi-step queries.
- Integrate the LLM tool with the Climate Data Store (CDS).
- Incorporate C3S-specific quality assessments from the Evaluation and Quality Control (EQC) Jupyter Book into the LLM training pipeline to enable context-aware, scientifically grounded responses on dataset fitness for purpose.
- Integrate relevant scientific literature, including intercomparison studies and peer-reviewed publications, to enrich the LLM's ability to guide users in selecting, interpreting, and applying climate datasets effectively.
- Implement a retrieval-based architecture that enables the tool to draw from existing, high-quality Copernicus documentation and metadata, without requiring full-scale LLM re-training.
- Implement features such as real-time insights, dynamic query handling, and data visualization.
- Develop the user interface for easy interaction with the LLM tool.
- Ensure cross-disciplinary collaboration and multilingual support.
- Ensure sensitive information relating to the Data Store Services (DSS) infrastructure and any hardcoded secrets are not leaked in the output.
- Include backend logging and analytics functionality to support the automated generation of user intelligence reports**,** capturing query trends, topic clusters, pain points, and interaction metrics.

Expected deliverables:

- *LLM Tool Prototype V1:* An early working version, including basic natural language querying, simple UI, and initial CDS metadata integration. Enables early usability testing and feedback collection.
- *Integration with CDS:* LLM tool fully integrated with the Copernicus Data Store (CDS) and capable of retrieving climate and weather data in response to natural language queries.
- *Curated Knowledge Base for LLM Integration:* A structured and annotated collection of C3S-relevant materials prepared for machine interpretation, including content sources, formatting approach, metadata schemas used, and links to authoritative documentation
- *Knowledge Base and Training Dataset Report:* A comprehensive report detailing the curated sources used to train and fine-tune the LLM tool, including C3S EQC quality assessments, intercomparison studies, and other relevant scientific literature.
- *User Interface (UI) Design and Prototype:* A user-friendly interface that allows users to interact with the LLM tool through natural language, including query input, results display, and visualizations.
- *Real-time Data Insights:* Implementation of features to offer real-time insights, trends, and data visualizations based on user queries.
- *Documentation for LLM Tool:* Technical documentation outlining the design, architecture, APIs, and how the LLM tool interfaces with the data sources.
- *User Stories and Test Cases:* A set of user stories and associated test cases to validate the LLM tool’s functionality and integration with the CDS.

Close collaboration with ECMWF and third parties will be essential throughout the contract to ensure alignment of requirements and seamless integration with existing content sources. The Tenderer is expected to propose a work methodology which is flexible to handle incoming requirements and adapt to external dependencies.

3.3 WP3: AI-Based Experimental Quality Assessment Development

Experimental work shall be performed in this work package to develop AI-based quality assessment models that evaluate quality attributes such as reliability, completeness, and accuracy of CDS datasets and assess overall data quality. The model shall be able to detect biases, resolution limitations, and gaps in datasets and provide context-specific assessments (e.g., for policy analysis, scientific research, or educational purposes). Data quality metrics shall be transformed into readable, understandable reports. These reports will be tailored to the user's expertise (e.g., novice, researcher, policymaker) and will summarise key points such as dataset accuracy, temporal and spatial resolution, and potential biases.

Generative AI could feed the quality assessments it generates directly into the LLM tool, enabling it to contextually respond to user queries about data quality, recommend datasets based on the specific quality needs of the user (e.g., "This dataset is best for global analysis, but not ideal for regional studies due to resolution limitations") and generate real-time quality evaluations as new datasets or models are queried by users.

By incorporating generative AI into the quality assessment process, the LLM tool can provide automated, dynamic, and highly detailed evaluations of climate datasets, making it not only a resource for understanding the data but also for improving data selection, quality control, and decision-making. This integration would allow the tool to cater to a broader range of users — from experts needing fine-grained analysis to policymakers who need understandable, actionable insights from data quality assessments.

The development of AI-generated quality assessments shall be led in collaboration with a senior climate scientist**,** who will co-design evaluation criteria and validate outputs for alignment with the EQC framework, published intercomparison studies, and scientific best practices. Their oversight ensures that automated assessments are interpretable, accurate, and trusted by expert users.

The successful Tenderer shall:

- Develop and test machine learning or generative AI methods for producing automated or semi-automated quality assessments of selected CDS datasets.
- Identify suitable input data sources**,** including observational references, metadata, and existing quality flags, to support AI-driven evaluations
- Develop methods for automated validation of climate and weather data.
- Design evaluation criteria and metrics to compare AI-generated assessments with existing manual or expert-led assessments for validation.
- Explore integration pathways for incorporating AI-generated quality assessments into the LLM tool and the broader C3S data ecosystem
- Implement a feedback loop to improve the AI quality assessment models.
- Conduct user feedback on how the quality assessment feature impacts data discovery and decision-making.

Expected deliverables:

- *Technical Approach for AI-Driven Quality Assessments:* A report outlining the proposed AI methodologies, input data sources, evaluation criteria, and implementation plan.

- *Prototype AI Model and Preliminary Results:* A working prototype of the AI model applied to a selected set of CDS datasets, accompanied by initial quality assessment outputs and comparison with existing evaluations.
- *Validation and Benchmarking Report:* Analysis of the performance of AI-generated assessments, including comparisons with expert-led assessments and documentation of accuracy, explainability, and user relevance.
- *Integration Feasibility and Recommendations:* A summary report assessing the feasibility of operational integration of AI-generated quality assessments into the LLM tool and C3S workflows, including technical and governance considerations

3.4 WP4: Testing, Rollout and User Interaction Analysis

This work package shall deal with testing the LLM tool and quality assessment features in real-world scenarios, evaluate the system's performance, and gather user feedback to refine the tool.

Testing and rollout will follow a phased approach. An internal pre-operational rollout will begin in Month 15, targeting expert users and stakeholders for structured testing under near-operational conditions. This will be followed by a public beta release in Month 21, offering all core functionalities to a broader user base for live feedback, usage metrics, and final refinements prior to operational transition.

In addition to structured usability testing, this work package will include the analysis of anonymised user interactions collected during the internal and public beta phases. This analysis will help identify common queries, data discovery patterns, points of confusion, and emerging themes. These insights will support both iterative tool improvements and strategic recommendations to ECMWF for future service development.

The successful Tenderer shall:

- Develop testing plans for functional, performance, and usability testing.
- Conduct user validation with diverse user groups (e.g., scientists, policymakers).
- Gather feedback on the LLM tool's usability**,** accuracy**,** and impact on data discovery.
- Perform stress testing and load testing to ensure scalability under large data and query volumes.
- Provide detailed evaluation reports and recommendations for improvement.
- Demonstrate no sensitive information (e.g. API keys, passwords, encryption keys, and certificates) is leaked in the output.
- Analyse anonymised user queries and interaction logs to extract patterns of potential user needs, identify common pain points and misunderstandings, and produce insights that support usability improvements and Copernicus service planning.

Throughout the testing and rollout phases, dedicated Level-2 user support shall be established and refined to ensure handling of user queries. During the internal rollout phase, user support workflows shall be ensured. Lastly, in the public beta phase, full support shall be available via the procedures described in section 4.4 which will require dedicated support by the successful Tenderer. These activities shall ensure a smooth user experience and prepare the support team for operational transition. Dedicated training materials shall be delivered as part of the handover process.

Expected deliverables:

- *Test Plan:* A comprehensive test plan covering functionality, usability, performance, integration, and user roles. Defines test cases, environments, acceptance criteria, and metrics.
- *Internal Testing Package:* Internal rollout release of the LLM tool to selected users including structured feedback mechanisms (Milestone M4.1).
- *Internal Testing and Usability Report*: A summary of results from internal testing, including issues identified, user feedback on functionality and usability, and initial performance metrics.

- *Public Beta Launch Package:* Public beta release including updated user interface, CDS integration and documentation (Milestone M4.2).
- *User Testing Report:* Aggregated feedback from public beta testers. Covers user satisfaction, clarity of responses, usefulness of recommendations, and perceived trust.
- *Performance and Scalability Report:* Final report on performance benchmarking (load, query throughput, latency) and scalability results. Includes stress testing outcomes.
- *Bug and Issue Tracker Summary:* Cleaned and resolved bug/issue list with status updates and categorisation from both internal and external testing phases.
- *Test Case Execution Results:* Documentation of all test case results across integration, performance, and user acceptance testing. Pass/fail metrics included.
- *Usability Evaluation Report:* Finalised UX evaluation with recommendations for operational phase improvements. Includes insights on multilingual support, accessibility, and interaction flow.
- *First User Intelligence Report:* First analysis of user behaviour from beta interactions: trends, pain points, and service improvement opportunities.
- *Final User Intelligence Summary:* Final aggregated report with evolution over time and recommendations for service evolution.

3.5 WP5: Documentation, Handover and Transition

This work package shall provide comprehensive documentation, reports and a transition plan to ensure the tool can be effectively maintained and used by a wider audience. It comprises the final phase and shall build on lessons from the public beta period (WP4) and prepare for operational transition and handover. Final documentation, training materials, and a support plan shall be delivered in Months 24–30.

As part of the final documentation package, the deliverables shall include a dedicated section addressing legal and ethical considerations related to the tool's outputs. This shall cover liability disclaimers for AI-generated responses, particularly where the tool may be perceived as interpreting scientific data or offering guidance. It shall also include proper attribution and citation of all data sources**,** including CDS and any brokered datasets, to ensure transparency and compliance with Copernicus data policies. Particular attention shall be given to a smooth transition between the current CDS Virtual Assistant ("Knowledge Duck") and the new service.

The successful Tenderer shall:

- Document all technical aspects, including system architecture, AI model deployment, and integration with the Knowledge Duck.
- Create user guides and training materials for end-users and administrators.
- Create at least two internal training sessions for ECMWF staff, a video on the use of the service (in collaboration with ECMWF) and relevant contractors and contribute to three information sessions (organised together with ECMWF).
- Develop a sustainability plan for ongoing system maintenance and updates.
- Develop a transition strategy, inspired by well-known examples such as Google Gemini integrating with Google Search.
- Provide recommendations for future iterations of the LLM tool, including potential scalability and feature expansion.
- Ensure smooth handover and knowledge transfer to relevant stakeholders.

Expected deliverables:

- *Operational Readiness Report:* Consolidated summary of testing outcomes, system performance, and lessons learned from the beta phase. Confirms system stability, usability, and maturity for operational transition.

- *Final System Documentation:* A comprehensive, technical document covering the LLM tool's architecture, integration details, AI-based quality assessments, and any customizations made during development. Prepared for handover to ECMWF.
- *User Manuals and Training Guides:* Detailed guides and tutorials for end-users and administrators on how to use the LLM tool, the Knowledge Duck integration, and quality assessment features.
- *Transition strategy:* Detailed plan on the integration of the CDS Virtual Assistant, providing a suite of options on how to integrate and improve existing functionalities.
- *Maintenance and Support Plan:* A plan outlining ongoing support, maintenance procedures, and recommendations for system updates and upgrades.
- *Recommendations for Future Development:* Strategic recommendations for scaling, expanding, and improving the LLM tool in future iterations, including possible new features, broader user integration, or advanced AI capabilities.

3.6 WP6: Scoping Future Extension to Other Copernicus Services

This work package shall explore the feasibility of extending the LLM-based tool's capabilities to address user queries across multiple Copernicus Services, beginning with CAMS. The work shall assess the interoperability of metadata between the CDS and ADS, identify gaps in EQC coverage and harmonisation, and propose potential integration points. The output will be a technical and strategic roadmap for the evolution of a future Services@ECMWF chat assistant**,** capable of supporting discovery and quality guidance across C3S, CAMS, and potentially other thematic services.

The successful Tenderer shall:

- Assess the feasibility of extending the LLM tool's capabilities to address user queries related to CAMS and the Atmosphere Data Store (ADS).
- Identify similarities and differences in metadata structures, access mechanisms, and EQC frameworks between C3S and CAMS.
- Engage with relevant CAMS and ECMWF stakeholders to understand user needs and operational constraints.
- Review user queries (e.g. GHG emissions, atmospheric composition) that overlap between C3S and CAMS and evaluate how the tool could provide context-aware cross-service recommendations.
- Propose technical options for enabling basic interoperability, including metadata mapping, redirect logic, or shared endpoints.

Expected deliverables:

- *Cross-Service Integration Feasibility Report:* A roadmap outlining metadata and EQC alignment challenges, technical options for integration, and recommended next steps toward a multi-service chatbot. Includes stakeholder feedback and user needs summary.

4 General Requirements

4.1 Implementation Schedule

ECMWF intends to award a single Framework Agreement for a period of maximum 30 months, which shall be implemented via a single multi-annual Service Contract expected to commence in Q1 2026.

Tenderers are expected to provide a detailed time plan and schedule as part of the tender response. The proposed time plan and schedule shall address the main tasks, inputs, outputs, intermediate review steps, milestones, deliverables and dates. Regular progress meetings will be held with ECMWF during the contract to assess project status, risks and actions.

4.2 WP0: Management and Coordination

This work package includes overall responsibility for day-to-day service management and coordination.

The following contract management aspects shall be considered and as needed briefly described in the proposal:

Plans for the mandatory reporting on implementation and forward planning.

Meetings (classified as tasks and listed in a separate table as part of the proposal):

- Organise quarterly teleconference meetings to discuss C3S service provision, service evolution and other topics, prepare corresponding summary minutes of these meetings and maintain a list of agreed actions and their status.
- ECMWF organises annual C3S General Assemblies. The Successful Tenderer is expected to attend these meetings with maximum 2 team members and contribute to discussions related to the topic of this ITT.
- Tenderers can propose additional contract internal meetings (e.g. kick-off meeting, regular meetings to monitor contract performance) as part of their response. Most such meetings should be held by remote participation.
- **Travel Prices:** Travel prices should be based on the European Commission's calculator [Table 3: Unit cost per distance band for air or combined air/rail travel, Commission Decision C(2024)5405], and consider a daily subsistence allowance not to exceed €300. Travel prices must reflect estimated actual costs and **must not include any profit margin**. If the proposed travel prices deviate from these reference values, a clear justification must be provided.

Quality assurance and control: the quality of reports and Deliverables shall be equivalent to the standard of peer-reviewed publications. The timely delivery as well as final quality check of the deliverables shall be ensured by the Successful Tenderer (in terms of content, use of ECMWF reporting templates for deliverables and reports (Microsoft Word), format, deliverable numbering and naming, typos…); all reports in this project shall be in English. Unless otherwise specified the specific contract Deliverables shall be made available to ECMWF in electronic format.

Communication management (incl. external and internal communication). Any external communication activity must be agreed with the ECMWF Copernicus Communication team in advance. This includes, but not exhaustively, communication planning, branding and visual style, media outreach, website and social media activity, externally facing text and graphical content and events. Agreed activity would also need to be evaluated and reported on once complete so that success measures and KPIs could be provided to the European Commission (cf. Clause 2.4.6 of the Framework Agreement).

Set of Key Performance Indicators (KPIs) suitable for monitor contract performance. The proposed KPIs shall be SMART (specific, measurable, actionable, realistic and time bound). The Successful Tenderer shall report to ECMWF on these KPIs as part of the Quarterly and Annual Implementation Reports. The proposed set of KPIs is expected to be updated regularly with ECMWF during the contract.

In particular, the KPIs shall focus on measuring the performance of the service. These shall include but are not limited to a) system performance (e.g. uptime rate, response time), b) search effectiveness (e.g. match rate of search results linking to reasonable CDS datasets and resources; supported by expert knowledge), and c) user satisfaction.

Risk Management: The proposal shall include a risk register that describes identified risks for each work package, along with a mitigation strategy for each of the identified risks. This mitigation strategy shall be composed by both preventive and corrective measures. The risk register shall be updated regularly by the Successful Tenderer, and any update (related to new risks, likelihood or impact) shall be reported during the progress review meeting, as well as part of the quarterly and annual implementation reports.

Particular attention shall be given to risks related to the high computational and storage demands of LLM-based systems. The risk register shall therefore address infrastructure resource risks. Tenderers shall provide an initial estimate of the expected infrastructure requirements and describe how these will be monitored and adjusted over time. Further risks to be covered relate to the integration of the LLM tool with existing services (e.g., potential conflict with the CDS Virtual Assistant Duck), FAIR compliance and data privacy.

Resources planning and tracking using the appropriate tools.

Subcontractor management, including conflict resolution, e.g. the prime contractor is responsible for settling disagreements, although advice/approval from ECMWF may be sought on the subject. A list of subcontractors describing their contribution and key personnel shall be provided, as well as backup names for all key positions in the contract. Tenderers shall describe how the Framework Agreement; in particular Clause 2.9 on Sub-contracting has been flowed down to all their subcontractors.

Management of personal data and how this meets the requirements of Clause 2.8 on Personal Data Protection and Annex 6 of the Framework Agreement.

List of minimum deliverables and milestones required as part of WP0, covering the contractual and financial reporting obligations towards ECMWF in line with the Terms and Conditions of the Framework Agreement (cf. Clause 2.3 and Annex 5):

WP0 Deliverables				
Deliverable#	*Responsible*	*Nature*	*Title*	*Due*
D0.y.z-YYYYQQ	Tenderer	Report	Quarterly Implementation Report QQ YYYY *QQ YYYY being the previous quarter*	Quarterly on, 15/04, 15/07 and 15/10
D0.y.z-YYYY	Tenderer	Report	Annual Implementation Report YYYY [Part 1] *YYYY being the Year n-1. Shall include:* • Quarterly Implementation Report Q4 YYYY ; YYYY *being the Year n-1* • *Preliminary financial information YYYY ; YYYY being the Year n-1*	Annually on 15/01
D0.y.z-YYYY	Tenderer	Report	Annual Implementation Report YYYY [Part 2] *YYYY being the Year n-1*	Annually on 28/02
D0.y.z-YYYY	Tenderer	Report	Annual Implementation plan YYYY *YYYY being the Year n+1*	Annually on 30/09
D0.y.z-YYYY	Tenderer	Other	Copy of prime contractor's general financial statements and audit report YYYY *YYYY being the Year n-1*	Annually around June (no associated cost)
D0.y.z	Tenderer	Report	Final report	At the end of the contract

Table 1: WP0 Deliverables

WP0 Milestones			
Milestone#	*Title*	*Means of verification*	*Due*
M0.1.1.MX	Progress Review meeting with ECMWF / Payment milestones	Minutes of meeting	At each Payment Milestone due date (~every 6 months)
M0.1.2.MX	Kick off meeting	Minutes of meeting	Month 1
M0.1.3.MX	Attendance to XXX meeting (e.g. General Assembly)	Minutes of meeting	Due one month after the meeting

Table 2: WP0 Milestones

4.3 Deliverables and Milestones

Deliverables should be consistent with the technical requirements specified in this document. A deliverable is a substantial, tangible or intangible good or service produced as a result of the contract. In other words, a

deliverable is an outcome produced in response to the specific objectives of the contract. Deliverables are subject to acceptance by the technical contract officers at ECMWF. All contract reports and documentation for this ITT shall be produced in English. The quality of reports and deliverables shall be equivalent to the standard of peer-reviewed publications and practice. Unless otherwise specified in the specific contract, deliverables shall be made available to ECMWF in electronic format (PDF/Microsoft Word/Microsoft Excel or HTML) via the Copernicus Deliverables Repository portal. The details will be agreed at the negotiation stage.

Each Deliverable shall have an associated resource allocation (person-months and financial budget). The total of these allocated resources shall amount to the requested budget associated with payroll.

Milestones should be designed as markers of demonstrable progress in service development and/or quality of service delivery, as applicable. They should not duplicate deliverables.

Tenderers shall complete the relevant table in Volume IIIA as part of their Tender, which includes the details of deliverables and milestones for all work packages and the schedules for each work package. Volume IIIA will be used by Tenderers to describe the complete list of deliverables, milestones, and schedules for each work package. All milestones and deliverables shall be numbered as indicated. All document deliverables shall be periodically updated and versioned as described in the tables.

ECMWF will provide the templates for reports and plans at the beginning of the contract. Reporting documents should be short and factual, following the guidance which will be provided by ECMWF during negotiations with the Successful Tenderer. Contract management and technical coordination is expected to amount to approx. 7-10% of the planned use of the resources.

4.4 Contribute to L2 support to Copernicus User Support Team

The objective of this task is to provide specialised support to users of the delivered products and services.

ECMWF has a well-established centralised User Support to provide multi-tiered technical support to all users of C3S data, products, tools and services. A service desk system is used for ticketing user requests and distributing these requests to specialists as needed. Dedicated staff at ECMWF promote and maintain self-help facilities (Copernicus Knowledge Base (CKB), user forum, FAQs and tutorials, etc.) and also provide individualised support on technical queries related to the CDS, data formats, data access, etc. In addition, ECMWF staff members provide specialised scientific support to address questions related to its industrial contributions to C3S, e.g., in the areas of global reanalysis and seasonal forecasting.

All C3S contractors are expected to contribute to the delivery of multi-tiered technical support for the data and/or services they provide. The Successful Tenderer shall provide expert (Level-2) support through a) the Jira ticketing system with agreed KPIs (for example, 85% of Level-2 tickets should be resolved within 15-business days), and/or b) the user forum[1] by monitoring topics and providing responses. The Successful Tenderer shall provide an email address which acts as the single contact point.

4.5 Communication

The Successful Tenderer shall support ECMWF in its communication activities for the C3S services, where they are related to the activities described in this ITT. Additional activities such as C3S website news items, C3S brochures and flyers, may be discussed on a case-by-case basis during the contract implementation.

4.6 Key Performance Indicators

The Successful Tenderer shall report to ECMWF on a set of Key Performance Indicators (KPIs) suitable for monitoring various aspects of service performance (by using the template included in Volume IIIB). The KPIs

[1] https://forum.ecmwf.int/

shall be designed to quantify various aspects of quality of service against the requirements described in this document. As part of the Tender, Tenderers shall specify a proposed set of KPIs appropriate for the service, e.g., relating to operational service delivery, quality, data access, user support, user satisfaction, etc., aligned with the requirements expressed above. These initial specifications shall be refined together with ECMWF during the first 6 months of the contract.

In particular, KPIs shall focus on measuring the performance of the service. These shall include but are not limited to a) system performance (e.g. uptime rate, response time), b) search effectiveness (e.g. match rate of search results linking to reasonable CDS datasets and resources; supported by expert knowledge), and c) user satisfaction.

4.7 Payment Plan

Tenderers can propose a Payment Plan in ITT Volume IIIA "Pricing and deliverables" (cf. Excel spreadsheet "Payment Plan preparation"):

- The Payment Milestones should relate to the deliverables and milestones delivered during the corresponding Payment Milestone period (e.g. the payment covering the period January-June would only relate to the deliverables and milestones whose due dates are part of the same period).
- The frequency of Progress Review Meetings might be adapted to synchronise with the anticipated date of completion of each Payment Milestone.
- In case of request for a payment at contract signature, please note that this should be duly substantiated (e.g. in terms of necessary investment prior to implementation or during first weeks/months for ensuring the initial set up of the project). It is necessary to relate this payment to activities subject to other Payment Milestones.

5 Tender Format and Content

General guidelines for the tender are described in Volume IIIB. This section describes specific requirements to prepare the proposal for this particular tender, along with guidelines for minimum content expected to be included in the proposal, additional to the content described in the general guidelines of Volume IIIB. This is not an exhaustive description and additional information may be necessary depending on the Tenderer's response.

5.1 Page Limits

As a guideline, it is expected that individual sections of the Tenderer's response do not exceed the page limits listed below. These are advisory limits and should be followed wherever possible, to avoid excessive or wordy responses.

Section	*Page Limit*
Executive Summary	2
Track Record	2 (for general) and 2 (per entity)
Quality of Resources to be Deployed	2 (excluding Table 1 in Volume IIIB and CVs with a maximum length of 2 pages each)
Technical Solution Proposed	20 (Table 2 in Volume IIIB, the section on references, publications, patents and any pre-existing IPR is excluded from the page limit and has no page limit)
Management and Implementation	6 (excluding Table 3, Table 5 and Table 6 in Volume IIIB) + 2 per each work package description (Table 4 in Volume IIIB)
Pricing Table	No limitation

Table 3: Page limits

5.2 Specific additional instructions for the Tenderer's response

The following is a guide to the minimum content expected to be included in each section, additional to the content described in the general guidelines of Volume IIIB. This is not an exhaustive description and additional information may be necessary depending on the Tenderer's response.

5.2.1 Executive Summary

The Tenderer shall provide an executive summary of the proposal, describing the objectives, team and service level.

5.2.2 Track Record

The Tenderer shall demonstrate for itself and for any proposed subcontractors that they have experience with relevant projects in the public or private sector at national or international level. ECMWF may ask for evidence of performance in the form of certificates issued or countersigned by the competent authority.

5.2.3 Quality of Resources to be Deployed

The Tenderer shall propose a team providing the skills required for providing operational services that meet the technical requirements set out in Section 3. The team shall include a Service Manager with at least five years of experience in management of large-scale projects. The Tenderer shall describe the experience of the Service Manager and the technical project team in performing activities related to the various aspects of this tender.

The proposed team shall include a dedicated climate science lead with deep expertise in climate data evaluation, C3S datasets, and the EQC framework. This expert will ensure that scientific integrity is maintained throughout system design, AI development, and validation. Their role includes co-reviewing LLM outputs, helping structure dataset recommendation logic, and liaising with external EQC contributors to ensure alignment with C3S standards.

5.2.4 Technical Solution Proposed

The Tenderer shall give a short background to the proposed solution to demonstrate understanding of that solution and of the C3S context. This section shall also include information on any other third-party suppliers that are used as part of the technical solution, and a statement of compliance for each requirement formulated throughout this document, describing how the proposed solution maps to the requirements.

5.2.5 Management and Implementation

As part of the general project management description, and in addition to the guidance provided in Volume IIIB, the Tenderer shall consider the elements described in section 4.2 above.

6 Additional Information

6.1 Acronyms

API	Application Programming Interface
AI	Artificial Intelligence
C3S	Copernicus Climate Change Service
CAMS	Copernicus Atmosphere Monitoring Service
CDS	Climate Data Store
ECMWF	European Centre for Medium-Range Weather Forecasts
EQC	Evaluation and Quality Control
FAIR	Findable, Accessible, Interoperable, and Reusable
ITT	Invitation To Tender

LLM	Large Language Model
NLP	Natural Language Processing
RAG	Retrieval-Augmented Generation
VA	Virtual Assistant