



Invitation to Tender

Destination Earth Initiative

Development and Implementation of a Climate and Weather Chatbot for Destination Earth

Volume II: Specification of Requirements

ITT Ref: DE_394
ISSUED BY: ECMWF Administration Department Procurement Section
Date: 24 May 2024
Version: Final

Table of Contents

1	Introduction.....	3
2	Context	3
3	Contract summary	5
4	Technical specification.....	6
4.1	Large language model (LLM)	6
4.2	Finetuning and performance	6
4.3	The web application	7
4.4	Relation to other DestinE activities	7
4.5	Quality of the documentation	7
4.6	Improper input and hallucinations	7
4.7	IPR.....	7
4.8	Operation and maintenance	8
5	General requirements	8
5.1	Implementation schedule.....	8
5.2	Meetings.....	8
5.3	Deliverables and milestones.....	9
5.3.1	Documents and reports	9
5.3.2	Other related DestinE activities	9
6	Tender format and content	10
6.1	Page limits	10
6.2	Specific additional instructions for the Tenderer’s response.....	10
6.2.1	Executive summary.....	10
6.2.2	Track record	10
6.2.3	Quality of resources to be deployed	10
6.2.4	Technical solution proposed.....	11
6.2.5	Management and implementation plan.....	11
6.2.6	Key performance indicators.....	12
6.2.7	Diversity and inclusion.....	12

1 Introduction

Destination Earth (DestinE) is an initiative of the European Commission under the EU Digital Europe programme. By pushing the limits of computing, weather and climate sciences, DestinE is a cornerstone of the European Commission's efforts to boost Europe's digital capabilities and the Green Deal actions on climate change. It aims at supporting climate change adaptation policies and decision-making for reducing the impacts of extreme events.

Destination Earth (DestinE) is deploying several high-resolution, thematic digital replicas (digital twins, DT) of the Earth system to monitor and simulate natural and human activities as well as their interactions. The DTs of DestinE will be used to enhance our understanding of the Earth system and to investigate how different weather and climate scenarios may impact the environment and, consequently, human life and societies.

The initial high-priority digital twins are implemented by the European Centre for Medium-Range Weather Forecasts (ECMWF) – one on climate change adaptation and one on weather-induced and geophysical extremes. These develop enhanced simulation systems, informed by observations, based on a new generation of Earth system models. These enhanced systems will not only allow to realistically represent the Earth system but will also produce information at precisely those scales where the impact of climate change and extremes are felt and where key processes are observed. They will thus allow users from impact-sectors to access and exploit such information for their specific application.

In an ambitious effort to make use of Machine Learning and Artificial Intelligence (ML/AI) to boost the DestinE information system and complement ongoing activities, this invitation to tender (ITT) is launched. The goal of this tender is to implement a climate and weather chatbot based on a large language model (LLM) to ease access to information for DestinE users and the wider community. As this technology is evolving very quickly, this tender document will not specify any specific solutions, but mostly requirements. Therefore, the first work package of this tender will be to conduct a business analysis to establish the detailed requirements and the technologies that are available at the time as well as their respective cost. ECMWF and the successful tendered will agree which solution to implement for the remaining part of the work.

The core outcome of this ITT should be a web-based chatbot that will converse with users and answer their specific queries, related to DestinE digital twins. The information source for the chatbot will also make use of ECMWF and Copernicus data and services.

2 Context

DestinE is funded by the European Union's Digital Europe programme and implemented through a partnership between the European Space Agency (ESA), the European Centre for Medium-Range Weather Forecasts (ECMWF) and the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT).

The first phase of DestinE, the implementation stage, covers the period 15 December 2021 – 14 June 2024. In this phase, the foundation infrastructure required to reach DestinE's ambitious goals is being configured and deployed and their capability is being demonstrated:

- The Core Service Platform (DESP; responsibility ESA) for providing many users with access to observations, simulations and models, evidence-based policy and decision-making tools, applications, and services, based on an open, flexible, scalable, evolvable, and secure cloud-based architecture.
- The Data Lake (DEDL; responsibility EUMETSAT) for handling the storage and access requirements for any input and output DestinE data that is offered to DestinE users via the seamless access through the DESP including near-data processing to maximize throughput and service scalability.

- The Digital Twin Engine (DTE; responsibility ECMWF) consisting of generic software infrastructures for workflows, extreme-scale simulation and data fusion, data handling and ML that allow exploiting the latest digital infrastructure technology for operating Earth-system DTs and their integration in the wider digital environment.
- The two high-priority DTs (responsibility ECMWF) for generating high-quality simulations and combining simulations and observations of the Earth system at unprecedented resolution to serve the EU's Green Deal policy priorities:
 - Weather-induced and geophysical extremes DT for providing capabilities for the assessment and prediction of environmental extremes at high spatial resolution and close to real-time decision-making support at continental, country, coastline, catchment, and city scales in response to meteorological, hydrological and air quality extremes.
 - Climate change adaptation DT for providing capabilities to support climate change adaptation policy and scenario testing at multi-decadal timescales aiming at a real breakthrough in resolution at regional and national levels.

The second phase of DestinE covers the period 15 June 2024 – 14 June 2026. This phase will focus on first operations, consolidation, maintenance, ramp-up, and continuous evolution of the System Components. DestinE users will thereby benefit from:

- Global Extremes DT simulations available at km-scale (2.8 to 4.4 km) on a timescale of 4 days ahead and event/impact-driven (on-demand) regional simulations at sub-km scale (500-750m) on a timescale of 2 days ahead
- Global climate projections up to 2050 at about 5km resolution, produced on a regular basis
- DTE performance and stability improvements, improving data access and data information management, while improving timeliness within the constraints of the available EuroHPC resources
- Responsive evolution of DESP services, co-designed with DESP users
- Growing portfolio of user-provided services on DESP
- Increasing maturity and fitness-for-purpose of the DestinE data portfolio and improved performance of federated access
- Further development of the uncertainty quantification approach, particularly through novel ML/AI techniques
- Further demonstration of co-developed use cases of DestinE capabilities in different impact sectors with selected users
- Developments towards a ML foundation model for the Earth system

Subsequent phases of DestinE will focus on full operationalisation of ongoing activities and improving Destination Earth content with a long-term objective of the availability of Destination Earth services.

The European Union has defined its approach to AI systems as one that balances opportunities and mitigates risks. Following the publication of an AI strategy for Europe in 2018¹ that already defined key principles for an ethical and legal framework of developing and using AI in Europe, the EU is one of the first world regions to implement comprehensive regulation of AI via the “AI Act”², a core element of the European Commission’s AI package presented in 2021. DestinE must be developed fully in line with the EU approach to AI and respect all ethical guidelines and principles as well as any legal constraints. Considering the specific context of DestinE, i.e., Earth System and impact sector modelling, some of these guidelines are more applicable than others. The work ECMWF intends to contract via this ITT should help implement the AI developments under DestinE with optimal coherence with European technological capabilities, opportunities, and values.

¹ COM(2018) 237 final: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52018DC0237>

² <https://www.europarl.europa.eu/legislative-train/theme-a-europe-fit-for-the-digital-age/file-regulation-on-artificial-intelligence>

3 Contract summary

The contract's main deliverable is a weather and climate web-based chatbot that needs to cover a wide variety of use-cases. The individual components should consist of:

- A Large Language Model (LLM)
- A system for requests to the LLM, (e.g., Retrieval-Augmented Generation RAG)
- A system to create, execute and process relevant data requests (e.g., Reasoning and Acting ReAct)
- A web interface

When referring to a "chatbot" in this tender it is implicitly including all of the above components.

The development of the chatbot will be structured through use cases. The use cases that the chatbot should be able to handle can broadly be split into two categories: documentation and support:

- Provide information about DestinE, that will cover scientific and technical documentation as well as policies. Where possible, synergies should be sought with Copernicus and ECMWF activities as well as policy and licencing matters.
- Provide help and support products and services, including example of software code when applicable.

The chatbot will have to be able to access a variety of data and documents, supporting the various use cases:

- Data and graphics produced by DestinE
- News and articles from the Destination Earth Joint Website, ECMWF's corporate website and DestinE dedicated website, newsletters
- Official documents, such as IPCC reports
- Scientific publications and technical memorandum
- Documentations, FAQs related to the various tools and services offered by DestinE
- User forums and support tickets

Documents will be in the following formats: Word, PDF, HTML and Markdown.

ECMWF, DestinE and Copernicus offer various APIs to retrieve past, present and future weather and climate information. These APIs can return graphical products (e.g., maps and time series as images) as well as numerical results (e.g., numerical values of a meteorological parameter at a point on the globe). The chatbot should be able to call these APIs when answering questions requires it. The European Union has 27 member states, and having the chatbot understanding and replying to queries in the language of the user would be a distinct advantage.

The chatbot will provide access to DestinE services and products, more specifically:

- An interface to existing DestinE APIs to retrieve graphical or numerical products.
- Be able to answer queries on questions that requires information from the above.

The first deliverable of the contract will be a business analysis where the Successful Tenderer is expected to identify any missing APIs that would help the implementation of the above feature. . The selection of the material (documentation and APIs) to be used to create the chatbot will be part of this first deliverable.

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

- ECMWF's experimental AI-based assistant³
- Automated Text Weather Forecast⁴

³ <https://chat.ecmwf.int>

⁴ <https://text-weather-new-neilgordon.replit.app/>

- Climate Foresight⁵

Tenderers should submit a proposal that addresses the areas outlined in the technical specification (Section 4) below. If a group of suppliers is involved in the delivery of the proposal, a prime contractor should be identified to lead the bid.

The technical selection criteria for the DestinE chatbot focus on demonstrated ML and AI expertise, awareness and understanding of relevant large language models and how they are pre-trained and fine-tuned, and how the results are disseminated. The technical requirements are described in the following section.

4 Technical specification

The Tenderer shall develop the requirements and implementation of a climate and weather chatbot that covers all of the below described features and tools.

4.1 Large language model (LLM)

Tenderers are free to propose solutions based on open-source version Large Language Model (LLM) such as Llama or Mistral. The preferred choice would be an LLM based in the European domain. The models should be accessible via an API, hosted in cloud service provider, or within the DESP infrastructure. In the latter case, the Tenderer is invited to specify the size of the resource required (RAM, CPU, GPU, storage, etc). For a costing model we refer to <https://www.ovhcloud.com/en-ie/public-cloud/prices/>.

The issue of GDPR must be considered. If a solution is based on finetuning an existing model, the cost and frequency of this activity must also be considered in the cost. Specifically, the training of ECMWF staff for further fine-tuning of the model needs to be considered.

4.2 Finetuning and performance

The performance of the finished system will be evaluated using several criteria, including:

- The relevance of the answer
- The speed at which the system provides an answer

It is expected that the first kind of use cases (e.g., questions about documentation) will be implemented using prompting techniques such as Retrieval-Augmented Generation (RAG). There are many improvements to RAG that lead to better results. Unfortunately, some of them add such a delay to the response that risks rendering the system unusable for interactive use. Most responses must be answered in under 10 seconds, with some exceptions, for which the time may be longer, provided there is meaningful visual feedback on the web application while the user is waiting.

It is expected that the second kind of use case (e.g., access to products via APIs) will be implemented using prompting techniques such as Reasoning and Acting (ReAct), that will invoke several tools, in order to provide the user with an answer. In that case ECMWF will be responsible for the performance of its own tools that run on ECMWF's internal cloud infrastructure or on DESP infrastructure and will work with the Successful Tenderer to ensure swift responses.

Any wait time should be accompanied by visual feedback to the users.

The relevance and accuracy of the answers will be judged by an expert panel using a set of pre-determined questions. The questions will be proposed by the Successful Tenderer in the business analysis. The final set of questions will be agreed between ECMWF and the Successful Tenderer. The answers will be evaluated

⁵ <https://www.nature.com/articles/s43247-023-01199-1>

against a baseline or general chatbot which has not been pre-trained and finetuned on the material. For the other evaluation criteria, please see section 6.2.6 for the KPI.

4.3 The web application

The chatbot will be made available as a web application. Users will start interacting with the system by typing a query in a text field. Answers should be accompanied by a series of web links to the relevant documents when available. Graphical answers must be presented inline as part of the conversation. Numerical answers must be presented to the users in ways that are easy to understand (tables, graphs). The system should record every question and answer for further tuning in an anonymised way. The user must be presented with a widget to rate the quality and relevance of the reply, which will also be recorded for further investigation. The collected feedback needs to be collected for future improvement of the chatbot.

The cost for hosting the web application on the DESP needs to be costed in the ITT. For a costing model we refer to <https://www.ovhcloud.com/en-ie/public-cloud/prices/>.

4.4 Relation to other DestinE activities

The Successful Tenderer shall make sure to liaise and collaborate with the activities carried out in the DE_340 contract procured by ECMWF where a chatbot is to be developed in the context of the Climate Change Adaptation Digital Twin activities. The proposed chatbot in DE_340 builds on the Climsight prototype and will be using an LLM to deliver local data from the Climate DT at local level using retrieval-augmented generation, fine-tuning these models with climate specific knowledge. The activities should complement each other and build synergies wherever possible and avoid duplication. The collaboration with DE_340 will need to be established in the first phase of the project and be included in the Business Analysis (see section 5.1 below).

The Successful Tenderer should also make sure that the contents of the chatbot take into consideration the output of the contract 'DE_398 Development and Implementation of Ethical Machine Learning Strategies for Destination Earth', currently being procured by ECMWF, which will deliver white papers on ethical and legal aspects of ML.

4.5 Quality of the documentation

RAG and fine-tuning are very sensitive to the quality of the documents, so part of the work will be to find a way to classify or rank documents based on a level of quality metric to be defined by the Successful Tenderer. This metric will be used to decide which documents should be used in finetuning or for RAG. This metric is expected to be a combination of an automated process (possibly using an LLM itself) and a series of criteria defined by the administrator of the system (e.g., source of the document). Tenderers should suggest how the quality of the metric is to be defined in their response.

4.6 Improper input and hallucinations

The chatbot should terminate any conversation that is not relevant to the business of DestinE, ECMWF or Copernicus (political, improper language, etc). Tenderers should describe how they plan to deal with model hallucinations in their response.

4.7 IPR

It is a condition of EU funding for DestinE that ownership of any Deliverables (as defined in Volume V Agreement) developed with DestinE funding passes from the suppliers to the European Union via ECMWF. Ownership will pass from the date of creation. All Background IPR (e.g., software and products) used by the Successful Tenderer to produce the results (Deliverables) will remain the property of the owner, e.g., the

Successful Tenderer. The Successful Tenderer will have to provide a royalty-free, non-exclusive, irrevocable, worldwide, and perpetual licence to Background IPR to the EU via ECMWF under the conditions set out in Volume V Agreement.

Developments or modifications to Background IPR which constitute Deliverables or Improvements and are created specifically for DestinE purposes will be owned by the European Union via ECMWF. A licence will be granted to the supplier to use the Deliverables for the provision of services. Upon request, suppliers may be granted a non-exclusive licence, at the discretion of ECMWF and subject to the approval by the European Commission, to use for other purposes the Deliverables which they have provided to DestinE.

4.8 Operation and maintenance

The system should always provide up-to-date information. Therefore, it should be possible for an administrator to add and remove sources of information. These sources are expected to be mostly URLs of websites to be automatically scrapped on a regular basis. Please note that the administrator will be an employee of ECMWF. Tenderers should provide a plan for knowledge transfer of how to update the chatbot in their response.

Tenderers should propose and include a period of support and maintenance of the software at ECMWF for a period of time after the delivering the chatbot. Tenderers should also provide a pricing mechanism for developments in the maintenance phase, see section 5.1 below for further details.

5 General requirements

5.1 Implementation schedule

ECMWF intends to award a contract, with a maximum value of €500k for the requested web-based chatbot along with documentation on how to update the content. The maximum duration of the contract is foreseen to be 24 months; the development phase of the chatbot is 12 months from the start of the contract, while the final 12 months of the contract are for maintenance and support, as well as potential further developments. The contract is expected to commence by Q4 2024.

In the development phase, the contractor is expected to deliver a “Business Analysis” within three months of signing the contract, where the requirements and scope of the chatbot is presented. The contractor is expected to deliver a beta version of the chatbot after 9 months from start of the contract, and the first operational version after 12 months from start of the contract.

Tenderers should include a period of running and maintaining of the software once the chatbot is operational. This should include costing for maintenance set at 2 days/month. The maintenance shall consist of fixing bugs. ECMWF shall be able to update and train the chatbot using the documentation provided by the Successful Tenderer. Tenderers should provide a costing model for expansion and development of the service. Depending on evolving needs and requirements, ECMWF may order such additional developments during or beyond the contract period.

Tenderers are expected to provide a detailed schedule as part of the tender response. The proposed time plan and schedule shall address the main tasks, inputs, outputs, intermediate review steps, milestones, and deliverables.

5.2 Meetings

Monthly progress meetings will be held (video conferencing) with ECMWF during the contract to assess contract status, risks, and actions and other topics that cut across various aspects of DestinE. ECMWF will organise annual physical meetings (in Bonn, occurring around April each year) to bring together all DestinE

capability providers, at which the Successful Tenderer is expected to be present. The cost of attending the physical meetings shall be covered by each Successful Tenderer and must be included in the tendered price. ECMWF may adjust meeting frequency as needed with the option of physical meetings at ECMWF’s Bonn duty station during the contract to demonstrate progress on this contract (up to 3 meetings in total). In addition, the Successful Tenderer may be asked to demonstrate or present their work at conferences and workshops on behalf of ECMWF and should allocate budget accordingly (2 conferences/ workshops).

In addition, the Successful Tenderer is expected to participate in a technical working group – which may also include other DestinE partners and relevant collaborators – aimed at discussing issues related to product definition, generation, and integration of the overall DestinE infrastructure. These discussions will be convened at regular intervals (quarterly) through videoconferencing.

5.3 Deliverables and milestones

Deliverables are to be defined by the Tenderer based on the requirements outlined above but should include as a minimum those listed below. They can be in the form of software, documents or reports, datasets, and support to other related DestinE activities, or other as appropriate.

<i>Nature</i>	<i>Title</i>
Report	Business Analysis: Requirements and detailed scope and evaluation criteria of the chatbot (within 3 months after signing the contract)
Software	Climate and Weather chatbot, a beta version within 9 months of start of the contract
Software	Climate and Weather chatbot, a first operational version within 12 months of start of the contract
Document	Full documentation on the chatbot and manual on how to update content
Document	Final deliverable of the project including user engagement analysis of the chatbot.

Each deliverable must 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 as detailed in Volume IIIA of this ITT.

Milestones should be designed as markers of demonstrable progress in capability development or quality of capability delivery, as applicable. They should not duplicate deliverables but provide auditable evidence of progress and as such should be part of the proposal and not incur additional costs.

5.3.1 Documents and reports

All project reports must be produced in English. Unless otherwise specified in the specific contract, deliverable documents and reports shall be made available to ECMWF in electronic format (Microsoft Office or compatible), via the DestinE Deliverables Repository portal; the details will be agreed at the negotiation stage. Please refer to Clause 2.3 and the Annex 5 of the Volume V Agreement for details on Reporting Obligations.

5.3.2 Other related DestinE activities

The Successful Tenderer is required to support the wider DestinE activities, for example the DestinE partnership activities, communication, and training and outreach. Sufficient resources for covering these aspects shall be foreseen and included in the tender price. Outreach activities will be organised by ECMWF during the period of the contract. In such instances, the contractors will be approached by ECMWF for support on developing and delivering content. Similarly, DestinE will require contributions to training material on relevant topics from the contractor. This is expected to amount to 2-3 hours per month.

Contractors shall not establish their own brand for the selected projects but adopt and use DestinE and ECMWF branding. A communications package (including guidelines, logos, and templates) will be provided by ECMWF at the start of the contract.

6 Tender format and content

General guidelines for the tender are described in Volume IIIB. Specific requirements to prepare the proposal for this tender are described in the next sub-sections.

6.1 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. If additional material is referenced (for example supplementary documentation or URL links to online information), a summary should still be provided in the core response.

<i>Section</i>	<i>Page Limit</i>
<i>Executive Summary</i>	1
<i>Track Record</i>	2 (for general) and 1 (per entity)
<i>Quality of resources to be Deployed</i>	2 (excluding Table 1 in Volume IIIB and CVs with a maximum length of 1 page each)
<i>Technical Solution Proposed</i>	10 (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)
<i>Management and Implementation</i>	3 (excluding Table 4 and Table 5 in Volume IIIB) + 1 per each Work package description (Table 3 in Volume IIIB)
<i>Pricing Table</i>	No limitation

Table 1: Page limits

6.2 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.

6.2.1 Executive summary

The Tenderer shall provide an executive summary of the proposal, describing the objectives, team and summarising the proposed technical solution and capability demonstration.

6.2.2 Track record

The Tenderer shall demonstrate for themselves and for any proposed subcontractors that they have experience with relevant projects. ECMWF may ask for evidence of performance in the form of certificates issued or countersigned by the competent authority. The Tenderer must include a brief description for a minimum of two recent contracts to demonstrate their capacities for undertaking the tasks foreseen.

6.2.3 Quality of resources to be deployed

The Tenderer shall propose a team providing the skills required for developing and demonstrating the solutions complying with technical requirements set out in Section 4. The team shall include a dedicated project manager with experience in management of projects of comparable size. The Tenderer shall describe the experience of the project manager and the technical project team in performing activities related to the various aspects of this tender.

6.2.3.1 Existing capabilities

Tenderers should present information outlining the strength of their present capabilities. For a Tenderer tasked with developing a climate and weather chatbot in the context of the Destination Earth (DestinE) initiative, these essential capabilities expected are the following:

- **Expertise in ML/AI:** The Tenderer must possess deep knowledge in Machine Learning and Artificial Intelligence, especially in Large Language Models (LLMs). This expertise is crucial for producing the chatbot
- **Expertise in Web Applications:** The Tenderer must have expertise in web development to produce the chatbot web application. This includes documented experience in design and maintenance of web applications.
- **Collaborative and Coordination Skills:** Given the need to closely liaise with ECMWF's DestinE teams and other stakeholders, the Tenderer must excel in collaborative work environments. This includes strong coordination, project management skills, and the ability to maintain open lines of communication to ensure alignment and constructive interaction across different components of the project.
- **Innovative and Creative Approach:** The ability to think innovatively and creatively is crucial for developing, engaging web application. The Tenderer should be able to present complex information in an accessible and engaging manner and propose creative solutions to challenges.

6.2.4 Technical solution proposed

The Tenderer shall include a brief executive summary describing the overall proposed technical solution to demonstrate their understanding of the DestinE context and the specific requirements of this tender. This section shall also include information on other third-party suppliers or solutions that are proposed for delivering the technical solution.

6.2.5 Management and implementation plan

The Tenderer shall provide a detailed implementation plan of proposed activities for the duration of the contract. Deliverables should be consistent with the technical requirements specified in Section 4.

The Tenderer is requested to include management and implementation activities within a dedicated work package (WPO). The maximum number of milestones is not prescribed, but they should be designed as markers of demonstrable progress in capabilities development or quality of capability delivery to keep progress monitoring manageable.

Adjustments to the proposed implementation plan can be proposed by the Successful Tenderer, depending on the needs for the evolution of the technical solution, changed user requirements, or other requirements, but must be agreed to by ECMWF.

As part of the general project management description the Tenderer shall consider the following elements (this is not an exhaustive list):

- Annual work Plan and Semestrial, annual and final reports shall be provided in accordance with the Volume V Agreement Clause 2.3 and Annex 5.
- Monthly videoconferencing with ECMWF and a proposal for involvement of ECMWF in major project reviews shall be provided as part of the management plan. The Tenderer is responsible for the organisation of such meetings, including provision of minutes.
- If relevant, a list of sub-contractors and details of their contribution, key technical personnel involved in the contract, legal names and addresses shall be provided. The Tenderer shall describe how the Volume V Agreement, particularly Clause 2.9, has been communicated to all their sub-contractors.
- The Tenderer shall describe in the Proposal the management of personal data and how this meets the requirements of Clause 2.8 and Annex 6 of Volume V Agreement.

The table below provides the template to be used by the Tenderer to describe the complete list of deliverables, milestones, and schedules for this work package. All milestones and deliverables shall be numbered as indicated. All document deliverables shall be periodically updated and versioned as described in the table. Tenderers shall provide preliminary versions of the completed tables as part of their bid.

Deliverables for this work package shall include the following reports:

WPO Contractual Obligations Template			
<i>#</i>	<i>Nature</i>	<i>Title</i>	<i>Due</i>
D0.y.z-YYYY	Report	Semestrial Implementation Report (Jan-June 202X). This includes a specific Financial Report	15/07/202X
D0.y.z-YYYY	Report	Annual Implementation Report 2024 YYYY being the Year n-1 This includes a specific Financial Report	15/01/2025
D0.y.z	Report	Final Implementation Report	Within 60 days after end of contract
D0.y.z-YYYY	Report	12-month Work Plan YYYY being the Year n+1	within 14 days of contract signature, and on 31 st August each year thereafter
D0.y.z-YYYY	Other	Copy of prime contractor's general financial statements and audit report YYYY (YYYY being the Year n-1)	Annually (no-cost associated)

6.2.6 Key performance indicators

Contractors shall provide a set of Key Performance Indicators (KPIs) suitable for monitoring various aspect of service performance, including (but not limited to):

- **Accuracy of chatbot:** The chatbot shall be able to give accurate and relevant answers to a number of pre-determined questions.
- **Speed of chatbot:** The chatbot will need to provide correct answers within a specified time. The time for providing relevant answers will have to be short enough to keep users engaged,
- **Relevance of answers:** The answers will need to be relevant in the context they are given, providing added value to the user.
- **Engagement metrics for the weather and climate chatbot:** Reach and engagement levels (e.g., unique users, shares, likes, and comments) of the chatbot, assessing effectiveness in engaging a broad audience.
- **Technical Expertise and Innovation:** Evaluation of the innovative approaches and technical solutions proposed, based on expert reviews, the novelty of the solutions, and their alignment with project goals.
- **Quality of Deliverables:** Assessment of the quality of all deliverables, based on predefined quality criteria and stakeholder feedback.
- **Project Management Effectiveness:** Efficiency and effectiveness of project management, measured by meeting milestones and deadlines, budget adherence, and the management of the bidding group partners (if applicable).

The KPIs, to be defined by the Tenderer, are subject to review by ECMWF and may be updated if necessary.

6.2.7 Diversity and inclusion

If multiple bidders present equally qualified proposals (discrepancy lower than 1%), ECMWF will take into consideration the diversity and gender balance of each bidder's organisation as a tiebreaker when making

the final decision. We recognise that diversity and a collaborative environment are essential for advancing scientific discovery and innovation, and we are dedicated to creating a culture that encourages and supports the contributions of individuals from all backgrounds. These varied backgrounds are particularly influential in adopting a holistic view of ethical AI and representing European values. As part of this commitment, we encourage bids from companies who share our values and demonstrate a commitment to diversity and inclusion in their own organisations. We believe that working with suppliers who support our efforts to create a more inclusive and diverse community is key to achieving our goals and driving progress forward in all our areas of activities. Therefore, the Centre encourages all potential bidders to take these values into consideration when submitting proposals.