



Invitation to Tender

Destination Earth Initiative

Development and Implementation of Climate Emulator for Destination Earth

Volume II: Specification of Requirements

ITT Ref: DE_393
ISSUED BY: ECMWF Administration Department Procurement Section
Date: 8 May 2024
Version: Final

Table of Contents

1	Introduction.....	3
2	Context	4
3	Contract summary	5
4	Technical specification.....	5
4.1	Computing	8
4.2	Additional requirements for software and data	8
4.2.1	Software.....	8
4.2.2	Data transfer and data formats	9
4.2.3	Quality control, evaluation, support and documentation.....	9
4.3	Communication plan	10
4.4	IPR.....	10
5	General requirements	10
5.1	Implementation schedule.....	10
5.2	Meetings.....	10
5.3	Deliverables and milestones.....	11
5.3.1	Documents and reports	11
5.3.2	User support	11
5.3.3	Other related DestinE activities.....	11
6	Tender format and content	12
6.1	Page limits	12
6.2	Specific additional instructions for the Tenderer’s response.....	12
6.2.1	Executive summary.....	12
6.2.2	Track record	12
6.2.3	Quality of resources to be deployed	12
6.2.4	Technical solution proposed.....	13
6.2.5	Management and implementation plan.....	13
6.2.6	Key performance indicators.....	14
6.2.7	Diversity and inclusion.....	16

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 and climate change.

Destination Earth (DestinE) is deploying several high-resolution, thematic digital replicas (digital twins, DTs) 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.

The DestinE Climate Digital Twin is operationalising multi-decadal climate simulations at km-scale resolution. Current runs are now being produced with 5-10km resolution with different model configurations and for different historical and future periods, as well as control simulations. Additional scenario runs will become available during the second phase of DestinE until June 2026.

Running the DestinE Climate Digital Twin is computationally expensive. Furthermore, only a small fraction of the generated data can be stored due to the very large volume of the output. The goal of this ITT is to develop a machine learning-based emulator for the DestinE Climate Digital Twin that is trained on its output and can reproduce aspects of existing runs (e.g. a timeslice, certain variables) for a specified forcing at a small fraction of the compute budget. It should thereby reproduce climatologically relevant properties of the DestinE Climate Twin, for example match observed decadal trends for historical control runs or major modes of variability such as El Nino Southern Oscillation (ENSO), without the need to provide an exact reproduction of the DestinE Climate Digital Twin output. The emulator, furthermore, should be able to use the stored data of the DestinE Climate Twin as initial condition and be able to reproduce high-resolution output in space and time. The ITT has an optional component to study the capabilities of the machine learning model to interpolate between different forcing scenarios. Additional funding is available for this component.

We expect the tender to build on the rapid progress on machine learning-based simulations for global medium range weather forecasting in the last two years, for which state-of-the-art models now provide skill comparable to the best conventional models at a small fraction of energy consumption and compute requirements. In particular, the emulator should support a spatial resolution of 25 km in its final configuration and perform a full system time integration that allows the recovery of relevant fields from the DestinE Climate Twin. While first steps towards such climate emulators have already been done in the literature, the tender needs to address significant technical challenges such as long-term stable time integration and training with input from different physics-based simulations.

The results of the tender will likely help to inform the role machine learning will play in the third phase of Destination Earth to complement the DestinE Extreme and Climate DTs.

2 Context

DestinE is funded by the European Union's Digital Europe programme and implemented through a partnership between the European Centre for Medium-Range Weather Forecasts (ECMWF), the European Space Agency (ESA), 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 key elements of the DestinE system are developed and deployed and their capability is being demonstrated:

- 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 (Extremes DT).
 - 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 (Climate DT).
- 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 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 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 DT projections up to 2040 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.

3 Contract summary

The aim of this ITT is to develop a machine learning-based DestinE Climate Twin emulator in form of well-documented and runnable code package for both training and inference, pre-trained neural network models, as well as a documentation that details the capabilities and limitations of the emulator with respect to expected properties. The deliverables must be suitable to serve as basis for future developments as part of DestinE.

The proposed solutions shall contribute to the requirements definition of future DTs and the evolution of the first high-priority DTs beyond phase 2 of DestinE. The Successful Tenderer will thus be required by ECMWF to contribute to relevant workshops and requirements collection activities.

Tenderers should submit proposals that address the specific technical requirements in 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.

We expect the delivery of software and data products at technical readiness level (TRL) of 7 or higher, which implies that the entry points for this contract need to be based on developments with sufficient maturity to achieve TRL7 in the given contract period.

The technical selection criteria for this Destination Earth ML/AI project focus on demonstrated ML/AI expertise as well as expertise on climate modelling and evaluation and relevant project experience in national or EU-funded projects. Furthermore, an innovative yet feasible technical approach needs to be proposed that aligns with the requested funding.

4 Technical specification

The tender is to provide a machine learning-based emulator for DestinE Climate Twin simulations meeting the following criteria:

- i. stable model integration for 30 years or more;
- ii. work at approximately 25 km spatial resolution (equivalent to 0.25 degree for an equi-angular grid);

¹ Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions- Artificial Intelligence for Europe, COM(2018) 237 final, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM%3A2018%3A237%3AFIN>

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

- iii. provide model output that is physically consistent, e.g. has realistic spherical harmonics spectra even for long integration times and realistic covariances between variables;
- iv. match the statistical properties of the climate simulations used in training as much as possible, e.g., match observed decadal trends for historical control runs, provide a realistic representation of the seasonal cycle in all areas of the globe; realistic time-mean state for, e.g., temperature and rainfall, circulation metrics with systematic biases equal or smaller to the training data, realistic representation of the observed characteristics of major modes of variability, e.g. El Nino Southern Oscillation, Pacific decadal oscillation, North Atlantic oscillation, and Atlantic multi-decadal variability;
- v. continue a simulation from a DestinE Climate Twin restart file/model output and produce physical fields with high resolution in time and space (e.g. hourly and 25 km resolution in space) to interpolate between sparse output;
- vi. at inference, provide at least 1 simulation year per 6h wall clock time on a single state-of-the-art accelerator.

For this tendered work, we envision an end-to-end machine learning model, i.e. without a separation into the components of conventional climate simulations. In well justified cases, we will also consider a component-based approach. The tender aims at deterministic models although generative/probabilistic ones will be considered. The machine learning methodology should reflect the state-of-the-art for large-scale machine learning for weather and climate and, in particular, build on the recent progress on machine learning-based medium-range weather forecasting. Existing implementations, e.g. ECMWF's Artificial Intelligence/Integrated Forecasting System (AIFS) which will become open source in 2024, can be used as a basis; this is not a requirement and the tender will be evaluated on the overall scientific merit as well as the feasibility of the proposed solution. More generally, it is expected that the tender uses best practices and libraries that are publicly available whenever possible.

The developed software will be able to train from the DestinE Climate Twin simulations and it is envisioned that, at a later stage, it could be extended for online training in parallel with a simulation run to make use of all the available data (only a small fraction of which can be stored long-term). The software produced will be well-designed and implemented, and thoroughly documented. Trained models will be provided with the software and be easy to run by domain scientists (e.g. by integrating the developed model into ECMWF's ai-models, which has been developed for this purpose, which will enable seamless integration into the DestinE workflows). The scientific capabilities of the developed emulator will be evaluated with respect to the aforementioned criteria and strengths and deficiencies will be well documented. Tenderers should provide a strategy for this evaluation.

For model development, publicly available climate simulation datasets can be employed, e.g. from CMIP-6 runs. Final training should be performed on DestinE Climate Twin model output, potentially complemented by runs from the H2020 nextGEMS and HEP EERIE projects. The developed emulator should satisfy all target criteria for training from one conventional DestinE Climate Twin model. The ability to make use of multiple models would add merit to a tender. Available training data from the DestinE Climate Twin are standard atmospheric variables at hourly resolution on pressure levels, and ocean, sea ice, and land data at daily resolution. If additional data will be required for a proposed emulator, this should be specified as part of the proposal so that it could be considered for future runs of the DestinE Climate Twin.

The tender allows for a possible extension to develop interpolation capabilities between the forcing scenarios in the training data. Up to 300k Euros additional funding is available in this case. It is currently unclear if this is possible and with what development effort but to obtain the additional funding Tenderers need to provide a clear strategy for improving interpolation capabilities between forcing scenarios as well as a sound methodology to evaluate their success.

For all of the above topics, the following technical criteria are relevant in order to achieve (at least) technical readiness level 7:

1. Meet the scientific criteria of the emulator specified above, e.g. on reproducing the statistics of the DestinE Climate Digital Twin, and document the capabilities.
2. Provide a well-structured and documented software based on state-of-the-art practices and with clearly defined API interfaces. Also specify ancillary software (incl. licenses) that is used.
3. Provide a comprehensive list of the data used for training and inference with respect to field types, data frequency, data latency and data quality, and typical/average data file sizes.
4. Describe the digital technology infrastructure needed to perform machine learning and deploy the machine learning solution. The summary should include the technical specifications of the computing resources being used.
5. Develop an open software framework that will eventually allow running the ML/DL tools within the production workflow of DestinE in close collaboration with ECMWF. This should detail both training and inference modes of ML emulator and training the emulator online in parallel with DestinE Climate DT runs.
6. Carry out preparatory work to interface with DTE developments where applicable, e.g. with a software infrastructure framework for workflows, cloud infrastructure integration, data analytics and data handling, considering the portability of the offered solution and operated at scale across distributed European large-scale computing and data handling infrastructures.

In addition, the Tenderer is required to specify necessary data sources and data quantities for training and inference. Suggest where help from DestinE with data access is expected. Specify how the Tenderer's workflow scales with increasing volumes of data.

Tenderers are invited to propose further developments towards the future improvement of climate emulation capabilities after DestinE's phase 2:

1. Climate emulation capabilities that can be developed as extensions of the proposed tender.
2. Improvement of the user friendliness and scalability (in view of increasing data volumes and spatial resolution of DT data) for ML/DL software tools when interacting with DestinE data.
3. Which additional DestinE data should be stored in the data lakes for constructing valuable ML datasets.
4. Prospective for training on streamed data from the climate twin, and envisaged technical work required for implementation.
5. API developments supporting code maintenance and software sustainability.
6. Other innovative contributions to material for DestinE communication, marketing, outreach and education, as agreed with ECMWF.

To allow efficient negotiations and minimise the complexity of any future adjustments, each development activity should be formulated, as far as possible, as a self-contained module detailed in its own work package including deliverables and milestones, with clearly assigned responsibilities.

The following sections describe specific requirements from several perspectives: technical (in terms of what components shall be delivered), capability demonstration and quality assessment.

4.1 Computing

DestinE relies on a partnership with the EuroHPC Joint Undertaking (JU) for its computing resources. The EuroHPC JU allows the European Union and the EuroHPC JU participating countries to coordinate their efforts and pool their resources for reaching new levels of supercomputing in Europe. DestinE relies on the computing centres hosting these HPC systems (e.g. CSC, CINECA, BSC, MeLuxina). This is complemented by the emerging DestinE Data Lake and Core Service Platform cloud infrastructure.

The main allocations for the HPC resources necessary for the DT production will be provided by the EuroHPC JU, including computing resources for further data processing. The present commitment by the EuroHPC JU for activities of strategic importance for Europe (e.g. DestinE and similar activities) is 5% of the total node-hour budget available to the JU.

The proposal should include an estimate for the number of GPU hours required to implement the tender, including an overview of how it is used for the different deliverables. The tender should also include a detailed listing of the storage requirements for the project. The Tenderer shall describe in detail any other specific dependencies, e.g. on software used or data governance (availability, usability, short/medium storage and persistence, etc.) that are important to facilitate the proposed work.

From the EuroHPC systems an approximate quota of 100,000-200,000 node hours on GPU-equipped nodes could be made available for the tender from the overall DestinE compute budget. The Successful Tenderer will, furthermore, be eligible and is expected to apply for additional EuroHPC compute time, or national one if applicable. It is expected that applicants have in-house compute resources for model development.

4.2 Additional requirements for software and data

4.2.1 Software

The Tenderer should, where possible, rely on open-source solutions and build on established best practices. If a non-open-source solution is chosen, the Tenderer must justify this choice, provide details on the license and its restrictions as well as indicative costs for the license that should be included in the overall price. This should include whether the software requires an ongoing subscription or if there is a perpetual right to use within the license terms. ECMWF will use this indicative cost to calculate the total cost of ownership of the system beyond the end of contract. ECMWF reserves the right to reject the usage of any non-open source software.

Given the challenging scale of the DestinE computing and data handling tasks, the Tenderer shall provide information on relevant software components, development priorities, and a potential roadmap for implementing machine learning workflows based on the emulator that are portable across novel processor, memory and storage technologies, and consider the very large volumes of digital twin data used for training. Furthermore, the software infrastructure implementation shall be coordinated with ECMWF as part of the Digital Twin Engine (DTE) development that aims to implement and sustain a generic framework that will serve present and future DT developments and production.

ECMWF leads the development of the DTE functionality, which will include support for software management, continuous integration and code deployment, and benchmarking for portability and performance assessment. In particular, the DTE will provide high-performance data production and data access API, supporting DT data access and data governance tasks. The Tenderer shall interface with this API and add developments and components to the data flow pipeline (e.g. considering data processing stages for efficient ML training) and co-develop such functionalities as relevant for the input DT data to ML processes and develop innovative solutions for data processing as required (e.g. data selection, data compression, interfacing with ML toolkits).

Selected output from DTs and other data sources will be stored on the DEDL and may be used in the ML process.

4.2.2 Data transfer and data formats

There are different levels of data production and data access anticipated in DestinE. The preferred solution is for DTs to directly make use of an abstraction layer for DT data production and data streaming that is separate from the user data access provided by the DEDL or DESP. The implementation of this abstraction is the responsibility of ECMWF and used by the DTs (as part of the DTE³) connecting to the underlying high-performance data handling service. This service will make an interface available to the DEDL and the DESP for the subsequent dissemination and user data access process. The Successful Tenderer should interface for the DT data access through this API. The preferred DestinE internal data format for production is the WMO GRIB edition 2 format (GRIB2) and the developed emulator should support this for both in- and output when making forecast (inference). The Successful Tenderer is recommended to achieve this using an ai-models⁴ plugin. The Tenderer should describe in detail how the ingestion of DT data into the ML/DL pipeline is anticipated, including any requirements/flexibility on the data pre-processing, e.g. dependency on a particular input data structure and format and any needs for reformatting/copying data. The DestinE high-performance data access service will include limited options for data cube⁵ access, and sub-selected extraction of data. Beyond the GRIB2 data produced by the DestinE high-performance data handling service, any additional potentially unformatted and user supplied datasets (e.g. derived data, other data sources, pointwise or vertical profile observational data, image data) should be specified and an explicit provision made in the proposal for actions such as reading, reformatting, and/or pre-processing such data.

Any dependencies impacting the timely delivery of machine learning models, constraints or requirements for interactivity must be clearly described (e.g. by providing a list of input and output fields for specific machine learning tools as well as the dimension of the fields including the frequency and latency required, etc.). These aspects form part of the evaluation process.

4.2.3 Quality control, evaluation, support and documentation

Quality control procedures (including automated processes) shall be established that validate the quality and fitness of the developed software. The precise methods should be proposed by the Tenderer and will be agreed as part of the negotiations. As a minimum, the software development should follow procedures according to industry standards, e.g. provisions for version control, coherent and uniform code styles, code reviews, issue and bug tracking, branching and merge strategies, continuous unit and acceptance testing followed by continuous integration.

Furthermore, an evaluation pipeline is required that validates the scientific skill of the developed emulator based on the aforementioned criteria. The evaluation pipeline is part of the expected deliverables.

The ML workflow must be documented to allow ML users to understand the result, the contributing data sources, and to be informed about configuration and version changes of the ML pipeline and tools used. In the event of ECMWF detecting possible problems with the outputs, providers are expected to give timely support to ensure prompt and effective resolution. A draft of details of response procedure and time shall be provided as part of the Tender response.

The Tenderer is required to document the scope of security and information management aspects to be provided and on the assets to be protected, according to ISO 27001.

³ See <https://digital-twin-engine.readthedocs.io/en/latest/> for more details on the interfaces available

⁴ <https://github.com/ecmwf-lab/ai-models>

⁵ https://en.wikipedia.org/wiki/Data_cube

4.3 Communication plan

The Tenderer shall include in their bid a clear Communication Plan, which will be further developed and agreed during the contract. This should be aimed at effectively disseminating of the outcomes and achievements of the DestinE Climate Twin emulator. The Successful Tenderer shall deliver communication material to support effective communication targeting a broad audience, including stakeholders in the scientific community, policy makers, industry leaders, and the public, ensuring the wide-reaching impact and visibility of DestinE's pioneering work in this area. The Successful Tenderer is expected to collaborate closely with the ECMWF DestinE communication team to produce a high-quality, engaging content portfolio. This content, ranging from press releases and social media posts to white papers and video materials, should be tailored to these different audiences, ensuring the clarity and accessibility of complex information. Regular updates and coordination meetings should be scheduled to align communication strategies and share resources effectively.

4.4 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.

5 General requirements

5.1 Implementation schedule

ECMWF intends to award a contract, with an estimated value of €800k for the development of the machine learning-based emulator with an additional up to €300k for the optional development of scenario interpolation capabilities. The maximum duration of the contract is foreseen to be 24 months and is expected to commence by quarter 4 of 2024.

The Tenderer is 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/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. They can be in the form of software, documents or reports, datasets and support to other related DestinE activities. Note the requirements related to the delivery of software and data have been described above (see Section 4). The requirements for all other types are described in the following subsections.

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 and/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 Word/PDF/Microsoft Excel 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 User support

The Tenderer is expected to contribute to the delivery of technical support for the data and functionality they provide. Such technical support shall take the form of a direct response to individual queries from ECMWF as required, as well as potential contributions to FAQs, user guides and knowledge bases. The Tenderer must cost this as a separate task within the work package action defined in section 4.2.3.

5.3.3 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 short description for a minimum of two recent contracts to demonstrate their capacities for undertaking the tasks foreseen. As supporting documents for contracts, ECMWF may request statements issued by the clients and contact them.

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.

6.2.4 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.

6.2.4.1 Existing capabilities

Tenderers should present information outlining the strength of their present capabilities in the following form:

- A brief description of the technological heritage of the proposed components.
- Publications or internal documentation describing the technical capabilities / user guides to support the adoption of similar solutions.
- Technical examples, in a similar or relevant context to illustrate the expertise of the team to deliver the proposed activities.
- A brief description of experience with large HPC systems and supporting software stacks, and of developing and running complex workflows at scale.
- A brief description of the capacity to handle big data.

6.2.4.2 Provision of technical capabilities

Tenderers shall describe in detail their proposed technical approach to develop, implement and support the technical solution in compliance with the technical requirements laid out in Section 4. The description shall include information on how the proposed solution maps onto the requirements formulated throughout this document.

Tenderers shall describe their plan to maintain and update existing capabilities in the course of the contract. If any new, technical or scientific developments are considered necessary during this period, these shall also be described, for context, even if they may not qualify for funding under this contract. Where insufficient information does not allow the fully compliant description of a technical solution, such case shall be highlighted to facilitate discussions at the clarification stage.

6.2.4.3 Computing and data handling resources

Tenderers should provide information about the anticipated computer systems required to develop the capabilities. DT data will be produced on (different) EuroHPC platforms, with dedicated access to HPC computing resources. Requirements with respect to HPC and ancillary services in terms of cloud computing, processing and data handling, data formats, data storage and data transmission should be specified separately (see section 4.1).

If other computing or hardware resources (guaranteed in-kind, subject-to-proposal in-kind, or charged to this contract) are provided, the estimated associated computing and data handling cost needs to be accompanied by information on the nature of the anticipated resource and the elements included in this cost, in as much detail as possible.

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 number of milestones is not prescribed, but they should be designed as markers of demonstrable progress in capabilities development and/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):

- Semestrial, annual and final reports shall be provided in accordance with the Volume V Agreement Clause 2.3 and Annex 5.
- A work plan is expected to be agreed at contract negotiation stage for solution delivery within the contracted end date, with demonstratable progress to coincide with DTE Phase-2.
- Monthly video-conferencing 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 (and chairing) of such meetings, including provision of minutes and any actions lists.
- 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, in particular 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:

WP0 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	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 report to ECMWF on a set of Key Performance Indicators (KPIs) suitable for monitoring various aspect of service performance, including (but not limited to):

- Capability development
- Capability demonstration
- Output quality
- Technical performance
- User support

The KPIs, to be defined by the Tenderer, are subject to review by ECMWF and may be updated if necessary. The Tenderer should propose KPIs such that the overall DestinE KPIs (see below) can be reported upon where applicable.

KPI	Method	Unit	Applicable Phase
Partnerships			
Hydrology Applications	Agreement implementation	Number	I-IV
Energy Applications	Agreement implementation	Number	I-IV
Food Applications	Agreement implementation	Number	II-IV
Health Applications	Agreement implementation	Number	II-IV
Trans-continuum	Agreement implementation	Number	I-II
Others	Agreement implementation	Number	II-IV
Digital Twin capabilities			
Spatial resolution/coverage of monitoring and prediction DT Extremes datasets	Monitor system usage/evolution; define DT improvement over existing systems	km/area	I-IV
Spatial resolution/coverage of monitoring and prediction DT Climate datasets	Monitor system usage/evolution; define DT improvement over existing systems m	km/area	I-IV
Temporal availability/timeliness of DT output for continuous production mode	Monitor system usage/evolution; define DT improvement over existing system	hours	I-IV
Temporal availability/timeliness of DT output for on-demand production mode	Monitor system usage/evolution; define cost-benefit of DT set-up improvement over existing systems	hours	I-IV
Availability of decision-ready information derived from DT output	Monitor DT output uptake per impact sector	%	II-IV
Digital Twin data uptake			
Number of (service-level) applications using DT-Extremes output	Monitor the actively running applications on core platform using output (simulations, observations)	Number	II-IV
Number of (service-level) applications using DT-Climat e output	Monitor the actively running applications on core platform using output (simulations, observations)	Number	II-IV
Number of (service-level) applications using full-resolution, high-frequency output	Monitor the actively running applications on core platform using output (simulations, observations)	CPU Hours	II-IV
Number of (service-level) applications using critical-path output	Monitor the actively running applications on core platform using output (simulations, observations)	Number	II-IV

Number of applications producing candidate models to be added to DT Engine	Monitor and evaluate the number of applications suitable for integration	Number	II-IV
Number of new datasets created from DT output	Monitor the datasets made available from core platform relative to original portfolio	Number	II-IV
Computing and data handling performance			
Amount of data sets pushed into the Data Lake(s)	Monitor the overall data flow across bridges	Number, TB	I-IV
Ingestion, access & usage of Earth observation data	Monitor the overall data flow across bridges and uptake by DT Engine	Number, TB	I-IV
Actual HPC node allocation for continuous production mode	Monitor the node-hour allocation on test systems and EuroHPC platforms	Node-hours/day	I-IV
Actual HPC node allocation for on-demand product mode	Monitor the node-hour allocation on test systems and EuroHPC platforms	Node-hours/day	I-IV
Sustained vs peak performance	Assess application specific sustained performance on test systems and EuroHPC platforms	%	I-IV
Digital Twin Engine			
Extreme-scale software component uptake	Monitor number of models/data assimilation systems employing DTE modules	Number	II-IV

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.