

Invitation to Tender

Destination Earth Programme

Machine Learning for Earth system Digital Twins

Volume II

ITT Ref: DE_371
ISSUED BY: ECMWF Administration Department Procurement Section
Date: 06 April 2023
Version: Final




Table of Contents

1	Introduction.....	4
2	Context	4
3	Contract summary	5
4	Technical specification.....	6
4.1	Computing	7
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, support and documentation.....	9
4.2.4	IPR.....	9
5	General requirements	10
5.1	Implementation schedule.....	10
5.2	Meetings.....	10
5.3	Deliverables and milestones.....	10
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	11
6.1	Page limits	11
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.....	12
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 and to prevent environmental degradation. It aims at supporting climate change adaptation policies and decision-making for reducing the impacts of extreme events.

Destination Earth (DestinE) will deploy 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.

Two high-priority digital twins will be implemented by the European Centre for Medium-Range Weather Forecast (ECMWF) – one on climate change adaptation and one on weather-induced and geophysical extremes. These will 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.

Artificial Intelligence (AI) and, more precisely, Machine and Deep Learning (ML and DL) are important for DestinE on many different levels, e.g. the processing, quality assessment and selection of observations, efficiency enhancements via emulation, accelerator utilization and optimisation, as well as the processing, post-processing and information extraction of output data depending on user, application area and decision-making needs.

Very recent developments around the use of ML in medium-range forecasting highlight the ability to learn complex dynamics from earth-system data. Equivalent ML models cannot be used in DestinE as there is no training data to learn models at global km-scale. However, the ML techniques developed in pursuit of that objective have wider use around Earth-system forecasting and can be deployed to create ML-generated data that is tethered to high-resolution simulations. Specifically, this invitation to tender (ITT) requires developing machine learning applications to improve uncertainty quantification and interpolate in time, with the purpose of better capturing uncertainty. By tethering synthetic ML-generated data to high-resolution simulations one can exploit the strengths of both digital twin data and machine learning technology.

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 will be configured and deployed and their capability will be demonstrated:

- The Core Service Platform (DESP; responsibility ESA) for providing a large number of 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 very 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.

In the subsequent phases of DestinE, these building blocks should evolve to enhance capabilities, add thematic foci, ingest the latest scientific developments and observational information, and exploit the emerging digital infrastructure ecosystem being created by the Digital Europe programme in Europe. Whilst the developments are ultimately targeting using DT data, it should be emphasised that this information will not be available at the start of the contract and therefore it is encouraged that tenderers build proposals using existing data sources at the outset and describe the stages to introduce DT data as this becomes available. In particular, the proposal should state what DT data is required/desired.

3 Contract summary

The aim of this ITT is to demonstrate that ML/DL based methodologies can augment DestinE datasets and products from the Weather-induced and geophysical extremes DT.

Under this procurement, the DT concept shall be enhanced with data-driven models tethered to high resolution simulations to provide synthetic model data. Specifically, machine learning models will generate physically-realistic synthetic ensemble members and provide physically-consistent interpolation in time. Note the initial stage of developing these models will need to be demonstrated using existing data (e.g. from data from operational numerical weather predictions at ECMWF), while designing methodologies and workflows which can leverage data from the weather extremes DT once available. A two-stage approach is envisaged, with the introduction of DT data towards a later phase of the development timeframe, subject to this being available (tenderers should consider the DT data availability risk when describing the programme of work). Tenderers should outline how they plan to blend data from multiple resolutions to leverage both the large volume of existing data and high resolution DT data. The Tenderer is expected to demonstrate the technical ability to interact with the DT work and data flow with a longer-term view to provide a repeatable ML/DL training and inference pipeline and integrating into the wider data access, product generation and data processing pipelines.

As the DT work and data flows are being developed in parallel, the Tenderer is requested to consider these strands when designing and presenting their proposals. The response should explain options for work and dataflow interaction that best suit the tenderers proposed solution.

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 proposed solutions shall contribute to the requirements definition of future DTs and the evolution of the first high-priority DTs beyond phase 1. The successful Tenderer will thus be required by ECMWF to contribute to relevant workshops and requirements collection activities.

The complex nature of the entire DestinE data production and user interaction workflow and the dependencies on the evolving technical infrastructure – hardware and software – may require adjustments to development plans as necessary over the durations of the contract. Mid-contract review will be scheduled to allow the tenderer and ECMWF to jointly analyse progress and decide on the most effective use of the available resources to deliver the outputs of the contract. At this point, amendments to adjust the priorities, or repurposing developments, will be considered and may be agreed by ECMWF at its sole discretion.

Tenderers should submit proposals that address both topics outlined in the technical specification (Section 4) below. If a consortium of suppliers are involved in the delivery of the proposal, a prime contractor should be identified to lead the bid. The tenderer shall provide and demonstrate the various ML/DL based software components, with a strong preference for building on open-source capabilities that will eventually facilitate their incorporation into the overall DestinE workflow.

Technical criteria for selection will include the scientific and technical credibility, feasibility of the demonstration (incl TRL), proven technology readiness, output quality and the realism of integrating new components into the DT work and dataflows, and the DTE. The technical requirements are described in the following sections. Tenderers should note that ECMWF will also consider the Total Cost of Ownership beyond the contract period and ultimately make an award decision based on the best interests of the Centre taking value for money into account.

4 Technical specification

Proposals shall use machine learning to provide valuable synthetic earth system data tethered to high resolution simulations to further bolster the uncertainty representation and utility of such simulations.

Specifically:

- Use generative ML techniques to increase the ensemble size by creating additional physically-consistent ensemble members tethered to a small ensemble, or single member deterministic input. Additional ensemble members should be shown to improve the statistical properties of the existing ensemble, including the spread-error relationship, probabilistic skill scores and representation of extreme events.
- Tethering to model trajectories to provide higher temporal frequency data than the existing model output frequency, capable of reproducing hourly model data from six-hourly or coarser inputs. Such a tool can produce significant saving for data archiving and data dissemination whilst retaining high temporal frequency data useful for downstream applications, e.g. wind power forecasting. Well calibrated uncertainty information should be provided, for example through an ensemble of plausible trajectories. Success in this application will enable future DT runs to output data at coarser temporal resolution easing IO constraints.

For both tasks the following requirements should be met:

- Machine learning output should have coherent structure across space, time and variables. For example, physical consistency for the zonal and meridional direction of the wind vector as well as physical consistent covariances between variables.
- Output should have comparable energy distribution across spatial scales as the training data.
- The machine learning framework should be capable of training and operating at multiple spatial resolutions, with high predictive performance across all spatial resolutions seen during training. Specifically, the framework should be able to train and run inference with both 10km operational forecasting data and km-scale data from the digital twin simulations.
- Stability and consistency of model predictions is a necessity and needs to be thoroughly assessed.

The successful Tenderer selected under this procurement will have a proven track-record in ML/DL methodology development and applications in, for example, environmental modelling, observation and simulation-observation data fusion, complex data processing and information extraction, uncertainty quantification and/or accelerated computing. The proposed solutions are required to fully comply with the DestinE criteria as defined in this section.

The successful Tenderer will produce new high-impact ML/DL methods and software solutions that can work close to the data generation location of the DTs of DestinE and eventually become integrated in the DTE. The output shall be a well documented software package, covering both training and inference functionality, trained ML/DL models for achieving the above tasks, and an assessment of product quality with respect to the chosen reference output.

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

1. Define and provide API interfaces with a clear specification of ancillary software (incl. licenses) and data formats, the latter with respect to field types, data frequency, data latency and data quality, and typical/average data file sizes.
2. Specify necessary data sources and data quantities, what access is anticipated and the regularity of data requests. Suggest where DestinE could help with data access. Specify how your workflow scales with increasing volumes of data and substantially increased spatial resolution as expected from DTs.
3. 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.
4. 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/DL tools.
5. Demonstrate the impact of the developments that illustrates how the new methodologies contribute to enhanced output frequency at sufficient quality, timeliness and societal relevance in an end-to-end DT workflow. Output quality and timeliness should be documented and compared with leading existing methodologies.
6. Carry out preparatory work to interface with DTE developments where applicable, e.g. with a software infrastructure framework for workflows, cloud infrastructure integration, simulation and data assimilation codes, 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.

Tenderers are invited to propose further developments towards the future improvement of such capability provision after DestinE's phase 1:

1. 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.
2. Which additional DestinE data should be stored in the data lakes for constructing valuable ML datasets.
3. API developments supporting code maintenance and software sustainability.
4. 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). 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) as well as on computing centres that are organised through the partnership for advanced computing in Europe (PRACE) to create a sustainable, distributed yet well-connected European computing, data production and data exchange framework. This is complemented by the emerging DestinE DEDL and DESP cloud infrastructure once available.

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 and not to exceed 10% of the total allocation per system.

The Tenderer shall provide a detailed cost estimate for any required HPC, networking, cloud computing and storage, and quantify any in-kind contributions, if applicable. In general, the Tenderer is not expected to deliver any hardware. A usage profile should include the estimated resource consumption over the duration of the contract – from the development through to integration and testing, highlighting when resources will be needed to achieve project milestones. From the above systems an approximate quota of 10,000 node hours on GPU-equipped nodes could be made available. The tenderer should provide a budget for the number of node hours requested and outline additional hardware to be used. Inclusive in this budget should be resources to develop and demonstrate integration on LUMI and Leonardo. The Tenderer may outline proposals for exceedance of this quota with sufficient justification that this is required for successful completion of the objectives. The Tenderer shall also 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.

4.2 Additional requirements for software and data

4.2.1 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 roadmap for implementing machine learning workflows that achieve performance and are portable across novel processor, memory and storage technologies, and consider the largely increased volumes and spatial resolution of digital twin data. Furthermore, the software infrastructure implementation shall be coordinated with ECMWF as part of the DTE development that aims to implement and sustain a generic framework that will serve present and future DT developments and production. Notwithstanding the aforementioned, the Tenderer should, where possible, rely on open-source solutions. 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.

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).

EUMETSAT will host output from DTs and other data sources on the DEDL that may be used in the ML process. It is from the DEDL that the data will ultimately be made available to users via the ESA DESP. The

tenderer is required to incorporate this information as a design consideration in their proposed solution, specifying all data sources and associated access mechanism that are used for the purpose of this tender.

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 DTE) 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 Tenderer of this ITT 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). 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 cube¹ 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, support and documentation

Quality control procedures (including automated processes) shall be established that validate the quality of the ML outputs and the correct representation of DT data. The precise methods should be proposed by the Tenderer and will be agreed as part of the negotiations. As a minimum, the software development where applicable 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.

The ML workflow must be documented to allow 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.

4.2.4 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

¹ https://en.wikipedia.org/wiki/Data_cube

will have to provide a royalty-free, non-exclusive, irrevocable, worldwide and perpetual license 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.

Modifications to closed-source software should be factored out as software extensions that can be uniquely identified and will be (as Deliverables or Improvements) owned by the EU via ECMWF, with the intention of open sourcing the software extensions subject to approval from the European Commission. If the deliverables include Background IPR, such Background IPR will be licensed to the Union unless there is an agreement on the transfer of ownership to the European Union.

A license 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 ca.€700-800k, for a maximum duration of 24 months, expected to commence by October 2023.

The successful 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. A roadmap of future developments beyond the contracted period is also highly desirable.

5.2 Meetings

Regular progress meetings will be held (video conferencing) with ECMWF during the contract to assess contract status, risks, and actions. ECMWF will organise annual physical meetings to bring together all DestinE capability providers, at which the successful Tenderer is expected to be present with 2 persons. The successful Tenderer must attend monthly (video-conferencing) meetings to review progress and other topics that cut across different aspects of DestinE. 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. 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 through video-conferencing.

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.

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 particular 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 brief summary should still be provided in the core response.

<i>Section</i>	<i>Page Limit</i>
<i>Executive Summary</i>	2
<i>Track Record</i>	2 (for general) and 2 (per entity)
<i>Quality of resources to be Deployed</i>	2 (excluding Table 1 in Volume IIIB and CVs with a maximum length of 2 pages each)
<i>Technical Solution Proposed</i>	30 (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>	10 (excluding Table 4 and Table 5 in Volume IIIB) + 2 per each Work package description (Table 3 in Volume IIIB)
<i>Pricing Table</i>	No limitation

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, the Contracting authority. ECMWF may request statements issued by the clients and make contact with 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-1 review scheduled for April 2024,
- 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 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, 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:

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 2023 YYYY being the Year n-1 This includes a specific Financial Report	15/01/2024
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-Climate 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
Machine-learning software component uptake	Monitor number of models/data assimilation systems employing DTE modules	Number	II-IV
HPC efficiency gains in DT production	Monitor the change in time-to-solution/reduction of node allocations	%	II-IV
Data exploitation gains in DT production and use in applications	Monitor the change in data touched and actively used	%	II-IV

6.2.7 Diversity and inclusion

In the event that 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 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. 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.